

Phase Transition in Distance-Based Phylogeny Reconstruction*

Sébastien Roch[†]

August 31, 2011

Abstract

We introduce a new distance-based phylogeny reconstruction technique which provably achieves, at sufficiently short branch lengths, a logarithmic sequence-length requirement—improving significantly over previous polynomial bounds for distance-based methods and matching existing results for general methods. The technique is based on an averaging procedure that implicitly reconstructs ancestral sequences.

In the same token, we extend previous results on phase transitions in phylogeny reconstruction to general time-reversible models. More precisely, we show that in the so-called Kesten-Stigum zone (roughly, a region of the parameter space where ancestral sequences are well approximated by “linear combinations” of the observed sequences) sequences of length $O(\log n)$ suffice for reconstruction when branch lengths are discretized. Here n is the number of extant species.

Our results challenge, to some extent, the conventional wisdom that estimates of evolutionary distances alone carry significantly less information about phylogenies than full sequence datasets.

Keywords: Phylogenetics, distance-based methods, phase transitions, reconstruction problem.

*The results detailed here were announced without proof in [Roc08, Roc10].

[†]Department of Mathematics and Bioinformatics Program, University of California-Los Angeles, Los Angeles, California 90095, USA. Work supported by NSF grant DMS-1007144.

1 Introduction

The evolutionary history of a group of organisms is generally represented by a *phylogenetic tree* or *phylogeny* [Fel04, SS03]. The leaves of the tree represent the current species. Each branching indicates a speciation event. Many of the most popular techniques for reconstructing phylogenies from molecular data, e.g. UP-GMA, Neighbor-Joining, and BIO-NJ [SS63, SN87, Gas97], are examples of what are known as *distance-matrix methods*. The main advantage of these methods is their speed, which stems from a straightforward approach: 1) the estimation of a *distance matrix* from observed molecular sequences; and 2) the repeated agglomeration of the closest clusters of species. Each entry of the distance matrix is an estimate of the evolutionary distance between the corresponding pair of species, that is, roughly the time elapsed since their most recent common ancestor. This estimate is typically obtained by comparing aligned homologous DNA sequences extracted from the extant species—the basic insight being, the closer the species, the more similar their sequences. Most distance methods run in time polynomial in n , the number of leaves, and in k , the sequence length. This performance compares very favorably to that of the other two main classes of reconstruction methods, likelihood and parsimony methods, which are known to be computationally intractable [GF82, DS86, Day87, MV05, CT06, Roc06].

The question we address in this paper is the following: Is there a price to pay for this speed and simplicity? There are strong combinatorial [SHP88] and statistical [Fel04] reasons to believe that distance methods are not as accurate as more elaborate reconstruction techniques, notably maximum likelihood estimation (MLE). Indeed, in a typical instance of the phylogenetic reconstruction problem, we are given *aligned DNA sequences* $\{(\xi_l^i)_{i=1}^k\}_{l \in L}$, one sequence for each leaf $l \in L$, from which we seek to infer the phylogeny on L . Generally, all *sites* $(\xi_l^i)_{l \in L}$, for $i = 1, \dots, k$, are assumed to be independent and identically distributed according to a Markov model on a tree (see Section 1.1). For a subset $W \subseteq L$, we denote by μ_W the distribution of $(\xi_l^i)_{l \in W}$ under this model. Through their use of the distance matrix, distance methods reduce the data to *pairwise sequence correlations*, that is, they only use estimates of $\mu_2 = \{\mu_W : W \subseteq L, |W| = 2\}$. In doing so, they seemingly fail to take into account more subtle patterns in the data involving three or more species at a time. In contrast, MLE for example outputs a model that maximizes the *joint probability of all observed sequences*. We call methods that explicitly use the full dataset, such as MLE, *holistic methods*.

It is important to note that the issue is not one of *consistency*: when the sequence length tends to infinity, the estimate provided by distance methods—just like MLE—typically converges to the correct phylogeny. In particular, under mild assumptions, it suffices to know the pairwise site distributions μ_2 to recover the

topology of the phylogeny [CH91, Cha96]. Rather the question is: how fast is this convergence? Or more precisely, how should k scale as a function of n to guarantee a correct reconstruction with high probability? And are distance methods significantly slower to converge than holistic methods? Although we do not give a complete answer to these questions of practical interest here, we do provide strong evidence that some of the suspicions against distance methods are based on a simplistic view of the distance matrix. In particular, we open up the surprising possibility that distance methods actually exhibit optimal convergence rates.

Context. It is well-known that some of the most popular distance-matrix methods actually suffer from a prohibitive sequence-length requirement [Att99, LC06]. Nevertheless, over the past decade, much progress has been made in the design of fast-converging distance-matrix techniques, starting with the seminal work of Erdős et al. [ESSW99a]. The key insight behind the algorithm in [ESSW99a], often dubbed the Short Quartet Method (SQM), is that it discards long evolutionary distances, which are known to be statistically unreliable. The algorithm works by first building subtrees of small diameter and, in a second stage, putting the pieces back together. The SQM algorithm runs in polynomial time and guarantees the correct reconstruction with high probability of any phylogeny (modulo reasonable assumptions) when $k = \text{poly}(n)$. This is currently the best known convergence rate for distance methods. (See also [DMR06, DHJ⁺06, Mos07, GMS08, DMR09] for faster-converging algorithms involving *partial* reconstruction of the phylogeny.)

Although little is known about the sequence-length requirement of MLE [SS99, SS02], recent results of Mossel [Mos04], Daskalakis et al. [DMR06, DMR11], and Mihaescu et al. [MHR09] on a conjecture of Steel [Ste01] indicate that convergence rates as low as $k = O(\log n)$ can be achieved when the branch lengths are sufficiently short, using insights from statistical physics. We briefly describe these results.

As mentioned above, the classical model of DNA sequence evolution is a Markov model on a tree that is closely related to stochastic models used to study particle systems [Lig85, Geo88]. This type of model undergoes a phase transition that has been extensively studied in probability theory and statistical physics: at short branch lengths (in the binary symmetric case, up to 15% divergence *per edge*), in what is called the *reconstruction phase*, good estimates of the ancestral sequences can be obtained from the observed sequences; on the other hand, outside the reconstruction phase, very little information about ancestral states diffuses to the leaves. See e.g. [EKPS00] and references therein. The new algorithms in [Mos04, DMR06, DMR11, MHR09] exploit this phenomenon by alternately 1) reconstructing a few levels of the tree using distance-matrix techniques and 2) es-

timating distances between *internal* nodes by reconstructing ancestral sequences at the newly uncovered nodes. The overall algorithm is *not* distance-based, however, as the ancestral sequence reconstruction is performed using a complex function of the observed sequences named *recursive majority*. The rate $k = O(\log n)$ achieved by these algorithms is known to be necessary in general. Moreover, the slower rate $k = \text{poly}(n)$ is in fact necessary for all methods—distance-based or holistic—outside the reconstruction phase [Mos03]. In particular, note that distance methods are in some sense “optimal” *outside* the reconstruction phase by the results of [ESSW99a].

Beyond the oracle view of the distance matrix. It is an outstanding open problem to determine whether distance methods can achieve $k = O(\log n)$ in the reconstruction phase¹. From previous work on fast-converging distance methods, it is tempting to conjecture that $k = \text{poly}(n)$ is the best one can hope for. Indeed, all previous algorithms use the following “oracle view” of the distance matrix, as formalized by King et al. [KZZ03] and Mossel [Mos07]. As mentioned above, the reliability of distance estimates depends on the true evolutionary distances. From standard concentration inequalities, it follows that if leaves a and b are at distance $\tau(a, b)$, then the usual distance estimate $\hat{\tau}(a, b)$ (see Section 1.1) satisfies:

$$\text{if } \tau(a, b) < D + \varepsilon \text{ or } \hat{\tau}(a, b) < D + \varepsilon \text{ then } |\tau(a, b) - \hat{\tau}(a, b)| < \varepsilon, \quad (1)$$

for ε, D such that $k \propto (1 - e^{-\varepsilon})^{-2} e^{2D}$. Fix $\varepsilon > 0$ small and $k \ll \text{poly}(n)$. Let T be a complete binary tree with $\log_2 n$ levels. Imagine that the distance matrix is given by the following oracle: on input a pair of leaves (a, b) the oracle returns an estimate $\hat{\tau}(a, b)$ which satisfies (1). Now, notice that for any tree T' which is identical to T on the first $\log_2 n/2$ levels above the leaves, the oracle is allowed to return the same distance estimate as for T . That is, we cannot distinguish T and T' in this model unless $k = \text{poly}(n)$. (This argument can be made more formal along the lines of [KZZ03].)

What the oracle model ignores is that, under the assumption that the sequences are generated by a Markov model of evolution, the distance estimates

$$(\hat{\tau}(a, b))_{a, b \in [n]}$$

are in fact *correlated random variables*. More concretely, for leaves a, b, c, d , note that the joint distribution of $(\hat{\tau}(a, b), \hat{\tau}(c, d))$ depends in a nontrivial way on the joint site distribution μ_W at $W = \{a, b, c, d\}$. In other words, even though the distance matrix is—seemingly—only an estimate of the pairwise correlations μ_2 ,

¹Mike Steel offers a 100\$ reward for the solution of this problem.

it actually contains *some* information about all joint distributions. Note however that it is not immediately clear how to exploit this extra information or even how useful it could be.

As it turns out, the correlation structure of the distance matrix is in fact *very informative* at short branch lengths. More precisely, we introduce in this paper a new distance-based method with a convergence rate of $k = O(\log n)$ in the reconstruction phase (to be more accurate, in the so-called Kesten-Stigum phase; see below)—improving significantly over previous $\text{poly}(n)$ results. Note that the oracle model allows only the reconstruction of a $o(1)$ fraction of the levels in that case. Our new algorithm involves a distance averaging procedure that implicitly reconstructs ancestral sequences, thereby taking advantage of the phase transition discussed above. We also obtain the first results on Steel’s conjecture beyond the simple symmetric models studied by Daskalakis et al. [DMR06, DMR11, MHR09] (the so-called CFN and Jukes-Cantor models). In the next subsections, we introduce general definitions and state our results more formally. We also give an overview of the proof.

Further related work. For further related work on efficient phylogenetic tree reconstruction, see [ESSW99b, HNW99, CK01, Cs02].

1.1 Definitions

Phylogenies. We define phylogenies and evolutionary distances more formally.

Definition 1 (Phylogeny) A phylogeny is a rooted, edge-weighted, leaf-labeled tree $\mathcal{T} = (V, E, [n], \rho; \tau)$ where: V is the set of vertices; E is the set of edges; $L = [n] = \{0, \dots, n-1\}$ is the set of leaves; ρ is the root; $\tau : E \rightarrow (0, +\infty)$ is a positive edge weight function. We further assume that all internal nodes in \mathcal{T} have degree 3 except for the root ρ which has degree 2. We let \mathbb{Y}_n be the set of all such phylogenies on n leaves and we denote $\mathbb{Y} = \{\mathbb{Y}_n\}_{n \geq 1}$.

Definition 2 (Tree Metric) For two leaves $a, b \in [n]$, we denote by $\text{Path}(a, b)$ the set of edges on the unique path between a and b . A tree metric on a set $[n]$ is a positive function $d : [n] \times [n] \rightarrow (0, +\infty)$ such that there exists a tree $T = (V, E)$ with leaf set $[n]$ and an edge weight function $w : E \rightarrow (0, +\infty)$ satisfying the following: for all leaves $a, b \in [n]$

$$d(a, b) = \sum_{e \in \text{Path}(a, b)} w_e.$$

For convenience, we denote by $(\tau(a, b))_{a, b \in [n]}$ the tree metric corresponding to the phylogeny $\mathcal{T} = (V, E, [n], \rho; \tau)$. We extend $\tau(u, v)$ to all vertices $u, v \in V$ in the obvious way.

Example 1 (Homogeneous Tree) For an integer $h \geq 0$, we denote by $\mathcal{T}^{(h)} = (V^{(h)}, E^{(h)}, L^{(h)}, \rho^{(h)}; \tau)$ a rooted phylogeny where $T^{(h)}$ is the h -level complete binary tree with arbitrary edge weight function τ and $L^{(h)} = [2^h]$. For $0 \leq h' \leq h$, we let $L_{h'}^{(h)}$ be the vertices on level $h - h'$ (from the root). In particular, $L_0^{(h)} = L^{(h)}$ and $L_h^{(h)} = \{\rho^{(h)}\}$. We let $\mathbb{HY} = \{\mathbb{HY}_n\}_{n \geq 1}$ be the set of all phylogenies with homogeneous underlying trees.

Model of molecular sequence evolution. Phylogenies are reconstructed from molecular sequences extracted from the observed species. The standard model of evolution for such sequences is a Markov model on a tree (MMT).

Definition 3 (Markov Model on a Tree) Let Φ be a finite set of character states with $\varphi = |\Phi|$. Typically $\Phi = \{+1, -1\}$ or $\Phi = \{A, G, C, T\}$. Let $n \geq 1$ and let $T = (V, E, [n], \rho)$ be a rooted tree with leaves labeled in $[n]$. For each edge $e \in E$, we are given a $\varphi \times \varphi$ stochastic matrix $M^e = (M_{ij}^e)_{i, j \in \Phi}$, with fixed stationary distribution $\pi = (\pi_i)_{i \in \Phi}$. An MMT $(\{M^e\}_{e \in E}, T)$ associates a state σ_v in Φ to each vertex v in V as follows: pick a state for the root ρ according to π ; moving away from the root, choose a state for each vertex v independently according to the distribution $(M_{\sigma_u, j}^e)_{j \in \Phi}$, with $e = (u, v)$ where u is the parent of v .

The most common MMT used in phylogenetics is the so-called general time-reversible (GTR) model.

Definition 4 (GTR Model) Let Φ be a set of character states with $\varphi = |\Phi|$ and π be a distribution on Φ satisfying $\pi_i > 0$ for all $i \in \Phi$. For $n \geq 1$, let $\mathcal{T} = (V, E, [n], \rho; \tau)$ be a phylogeny. Let Q be a $\varphi \times \varphi$ rate matrix, that is, $Q_{ij} > 0$ for all $i \neq j$ and

$$\sum_{j \in \Phi} Q_{ij} = 0,$$

for all $i \in \Phi$. Assume Q is reversible with respect to π , that is,

$$\pi_i Q_{ij} = \pi_j Q_{ji},$$

for all $i, j \in \Phi$. The GTR model on \mathcal{T} with rate matrix Q is an MMT on $T = (V, E, [n], \rho)$ with transition matrices $M^e = e^{\tau_e Q}$, for all $e \in E$. By the reversibility assumption, Q has φ real eigenvalues

$$0 = \Lambda_1 > \Lambda_2 \geq \dots \geq \Lambda_\varphi.$$

We normalize Q by fixing $\Lambda_2 = -1$. We denote by \mathbb{Q}_φ the set of all such rate matrices. We let $\mathbb{G}_{n,\varphi} = \mathbb{Y}_n \otimes \mathbb{Q}_\varphi$ be the set of all φ -state GTR models on n leaves. We denote $\mathbb{G}_\varphi = \{\mathbb{G}_{n,\varphi}\}_{n \geq 1}$. We denote by ξ_W the vector of states on the vertices $W \subseteq V$. In particular, $\xi_{[n]}$ are the states at the leaves. We denote by $\mathcal{L}_{\mathcal{T},Q}$ the distribution of $\xi_{[n]}$.

GTR models include as special cases many popular models such as the CFN model.

Example 2 (CFN Model) The CFN model is the GTR model with $\varphi = 2$, $\pi = (1/2, 1/2)$, and

$$Q = Q^{\text{CFN}} \equiv \begin{pmatrix} -1/2 & 1/2 \\ 1/2 & -1/2 \end{pmatrix}.$$

Example 3 (Binary Asymmetric Channel) More generally, letting $\Phi = \{+, -\}$ and $\pi = (\pi_+, \pi_-)$, with $\pi_+, \pi_- > 0$, we can take

$$Q = \begin{pmatrix} -\pi_- & \pi_- \\ \pi_+ & -\pi_+ \end{pmatrix}.$$

Phylogenetic reconstruction. A standard assumption in molecular evolution is that each site in a sequence (DNA, protein, etc.) evolves *independently* according to a Markov model on a tree, such as the GTR model above. Because of the reversibility assumption, the root of the phylogeny cannot be identified and we reconstruct phylogenies up to their root.

Definition 5 (Phylogenetic Reconstruction Problem) Let $\tilde{\mathbb{Y}} = \{\tilde{\mathbb{Y}}_n\}_{n \geq 1}$ be a subset of phylogenies and $\tilde{\mathbb{Q}}_\varphi$ be a subset of rate matrices on φ states. Let $\mathcal{T} = (V, E, [n], \rho; \tau) \in \tilde{\mathbb{Y}}$. If $T = (V, E, [n], \rho)$ is the rooted tree underlying \mathcal{T} , we denote by $T_-[\mathcal{T}]$ the tree T where the root is removed: that is, we replace the two edges adjacent to the root by a single edge. We denote by \mathbb{T}_n the set of all leaf-labeled trees on n leaves with internal degrees 3 and we let $\mathbb{T} = \{\mathbb{T}_n\}_{n \geq 1}$. A phylogenetic reconstruction algorithm is a collection of maps $\mathcal{A} = \{\mathcal{A}_{n,k}\}_{n,k \geq 1}$ from sequences $(\xi_{[n]}^i)_{i=1}^k \in (\Phi^{[n]})^k$ to leaf-labeled trees $T \in \mathbb{T}_n$. We only consider algorithms \mathcal{A} computable in time polynomial in n and k . Let $k(n)$ be an increasing function of n . We say that \mathcal{A} solves the phylogenetic reconstruction problem on $\tilde{\mathbb{Y}} \otimes \tilde{\mathbb{Q}}_\varphi$ with sequence length $k = k(n)$ if for all $\delta > 0$, there is $n_0 \geq 1$ such that for all $n \geq n_0$, $\mathcal{T} \in \tilde{\mathbb{Y}}_n$, $Q \in \tilde{\mathbb{Q}}_\varphi$,

$$\mathbb{P} \left[\mathcal{A}_{n,k(n)} \left((\xi_{[n]}^i)_{i=1}^{k(n)} \right) = T_-[\mathcal{T}] \right] \geq 1 - \delta,$$

where $(\xi_{[n]}^i)_{i=1}^{k(n)}$ are i.i.d. samples from $\mathcal{L}_{\mathcal{T},Q}$.

An important result of this kind was given by Erdos et al. [ESSW99a].

Theorem 1 (Polynomial Reconstruction [ESSW99a]) *Let $0 < f \leq g < +\infty$ and denote by $\mathbb{Y}^{f,g}$ the set of all phylogenies $\mathcal{T} = (V, E, [n], \rho; \tau)$ satisfying $f \leq \tau_e \leq g, \forall e \in E$. Then, for all $\varphi \geq 2$ and all $0 < f \leq g < +\infty$, the phylogenetic reconstruction problem on $\mathbb{Y}^{f,g} \otimes \mathbb{Q}_\varphi$ can be solved with $k = \text{poly}(n)$.*

This result was recently improved by Daskalakis et al. [DMR06, DMR11] (see also [MHR09]) in the so-called Kesten-Stigum reconstruction phase, that is, when $g < \ln \sqrt{2}$.

Definition 6 (Δ -Branch Model) *Let $0 < \Delta \leq f \leq g < +\infty$ and denote by $\mathbb{Y}_\Delta^{f,g}$ the set of all phylogenies $\mathcal{T} = (V, E, [n], \rho; \tau)$ satisfying $f \leq \tau_e \leq g$ where τ_e is an integer multiple of Δ , for all $e \in E$. For $\varphi \geq 2$ and $Q \in \mathbb{Q}_\varphi$, we call $\mathbb{Y}_\Delta^{f,g} \otimes \{Q\}$ the Δ -Branch Model (Δ -BM).*

Let $g^* = \ln \sqrt{2}$.

Theorem 2 (Logarithmic Reconstruction [DMR06, DMR11, MHR09]) *For $0 < \Delta \leq f \leq g < g^*$, the phylogenetic reconstruction problem on $\mathbb{Y}_\Delta^{f,g} \otimes \{Q^{\text{CFN}}\}$ can be solved with $k = O(\log n)^2$.*

Distance methods. The proof of Theorem 1 uses *distance methods*, which we now define formally.

Definition 7 (Correlation Matrix) *Let Φ be a finite set with $\varphi \geq 2$. Let*

$$(\xi_a^i)_{i=1}^k, (\xi_b^i)_{i=1}^k \in \Phi^k$$

be the sequences at $a, b \in [n]$. For $v_1, v_2 \in \Phi$, we define the correlation matrix between a and b by

$$\hat{F}_{v_1 v_2}^{ab} = \frac{1}{k} \sum_{i=1}^k \mathbb{1}\{\xi_a^i = v_1, \xi_b^i = v_2\},$$

and $\hat{F}^{ab} = (\hat{F}_{v_1 v_2}^{ab})_{v_1, v_2 \in \Phi}$.

Definition 8 (Distance Method) *A phylogenetic reconstruction algorithm $\mathcal{A} = \{\mathcal{A}_{n,k}\}_{n,k \geq 1}$ is said to be distance-based if \mathcal{A} depends on the data $(\xi_{[n]}^i)_{i=1}^k \in (\Phi^{[n]})^k$ only through the correlation matrices $\{\hat{F}^{ab}\}_{a,b \in [n]}$.*

²The correct statement of this result appears in [DMR11]. Because of different conventions, our edge weights are scaled by a factor of 2 compared to those in [DMR11]. The dependence of k in Δ is Δ^{-2} .

The previous definition takes a very general view of distance-based methods: any method that uses only pairwise sequence comparisons. In practice, most distance-based approaches actually use a specific *distance estimator*, that is, a function of \hat{F}^{ab} that converges to $\tau(a, b)$ in probability as $n \rightarrow +\infty$. We give two classical examples below.

Example 4 (CFN Metric) *In the CFN case with state space $\Phi = \{+, -\}$, a standard distance estimator (up to a constant) is*

$$\mathcal{D}(\hat{F}) = -\ln \left(1 - 2(\hat{F}_{+-} + \hat{F}_{-+}) \right).$$

Example 5 (Log-Det Distance [BH87, Lak94, LSHP94, Ste94]) *More generally, a common distance estimator (up to scaling) is the so-called log-det distance*

$$\mathcal{D}(\hat{F}) = -\ln |\det \hat{F}|.$$

Loosely speaking, the log-det distance can be thought as a generalization of the CFN metric. We will use a different generalization of the CFN metric. See section 1.3.

1.2 Results

In our main result, we prove that phylogenies under GTR models of mutation can be inferred using a distance-based method from $k = O(\log n)$ sequence length.

Theorem 3 (Main Result) *For all $\varphi \geq 2$, $0 < \Delta \leq f \leq g < g^*$ and $Q \in \mathbb{Q}_\varphi$, there is a distance-based method solving the phylogenetic reconstruction problem on $\mathbb{Y}_{\Delta}^{f,g} \otimes \{Q\}$ with $k = O(\log n)$.³*

Note that this result is a substantial improvement over Theorem 1—at least, in a certain range of parameters—and that it matches the bound obtained in Theorem 2. The result is also novel in two ways over Theorem 2: only the distance matrix is used; the result applies to a larger class of mutation matrices. A weaker version of the result stated here was first reported without proof in [Roc08]. Note that in [Roc08] the result was stated without the discretization assumption which is in fact needed for the final step of the proof. This is further explained in Section 7.3 of [DMR11]. The new proofs presented here rely on recent joint work with Yuval Peres [PR11] on exponential moment bounds for quantities such as $\bar{\sigma}_a$.

In an attempt to keep the paper as self-contained as possible we first give a proof in the special case of homogeneous trees. This allows to keep the algorithmic details to a minimum. The proof appears in Section 3. We extend the result to general trees in Section 4. The more general result relies on a combinatorial algorithm of [DMR11].

³As in Theorem 2, the dependence of k in Δ is Δ^{-2} [Roc10].

1.3 Proof Overview

Distance averaging. The basic insight behind Steel’s conjecture is that the accurate reconstruction of ancestral sequences in the reconstruction phase can be harnessed to perform a better reconstruction of the phylogeny itself. For now, consider the CFN model with character space $\{+1, -1\}$ and assume that our phylogeny is homogeneous with uniform branch lengths ω . Generate k i.i.d. samples $(\sigma_V^i)_{i=1}^k$. Let a, b be two internal vertices on level $h - h' < h$ (from the root). Suppose we seek to estimate the distance between a and b . This estimation cannot be performed directly because the sequences at a and b are not known. However, we can try to *estimate* these internal sequences. Denote by A, B the leaf set below a and b respectively. An estimate of the sequence at a is the (properly normalized) “site-wise average” of the sequences at A

$$\bar{\sigma}_a^i = \frac{1}{|A|} \sum_{a' \in A} \frac{\sigma_{a'}^i}{e^{-\omega h'}}, \quad (2)$$

for $i = 1, \dots, k$, and similarly for b . It is not immediately clear how such a *site-wise* procedure involving *simultaneously* a large number of leaves can be performed using the more aggregated information in the correlation matrices $\{\hat{F}^{uv}\}_{u,v \in [n]}$. Nevertheless, note that the quantity we are ultimately interested in computing is the following estimate of the CFN metric between a and b

$$\bar{\tau}(a, b) = -\ln \left(\frac{1}{k} \sum_{i=1}^k \bar{\sigma}_a^i \bar{\sigma}_b^i \right).$$

Our results are based on the following observation:

$$\begin{aligned} \bar{\tau}(a, b) &= -\ln \left(\frac{1}{k} \sum_{i=1}^k \left(\frac{1}{|A|} \sum_{a' \in A} \frac{\sigma_{a'}^i}{e^{-\omega h'}} \right) \left(\frac{1}{|B|} \sum_{b' \in B} \frac{\sigma_{b'}^i}{e^{-\omega h'}} \right) \right) \\ &= -\ln \left(\frac{1}{|A||B|e^{-2\omega h'}} \sum_{a' \in A} \sum_{b' \in B} \left(\frac{1}{k} \sum_{i=1}^k \sigma_{a'}^i \sigma_{b'}^i \right) \right) \\ &= -\ln \left(\frac{1}{|A||B|e^{-2\omega h'}} \sum_{a' \in A} \sum_{b' \in B} e^{-\hat{\tau}(a', b')} \right), \end{aligned}$$

where note that the last line depends only on distance estimates $\hat{\tau}(a', b')$ between leaves a', b' in A, B . In other words, through this procedure, which we call *exponential averaging*, we perform an *implicit* ancestral sequence reconstruction using only distance estimates. One can also think of this as a variance reduction technique. When the branch lengths are not uniform, one needs to use a *weighted* version of (2). This requires the estimation of path lengths.

GTR models. In the case of GTR models, the standard log-det estimator does not lend itself well to the exponential averaging procedure described above. Instead, we use an estimator involving the right eigenvector ν corresponding to the second eigenvalue Λ_2 of Q . For $a, b \in [n]$, we consider the estimator

$$\hat{\tau}(a, b) = -\ln \left(\nu^\top \hat{F}^{ab} \nu \right). \quad (3)$$

This choice is justified by a generalization of (2) introduced in [MP03]. Note that ν may need to be estimated.

Concentration. There is a further complication in that to obtain results with high probability, one needs to show that $\bar{\tau}(a, b)$ is *highly concentrated*. However, one cannot directly apply standard concentration inequalities because $\bar{\sigma}_a$ is *not bounded*. Classical results on the reconstruction problem imply that the variance of $\bar{\sigma}_a$ is finite—which is not quite enough. To show concentration, we bound the moment generating function of $\bar{\sigma}_a$.

1.4 Organization

In Section 2, we provide a detailed account of the connection between ancestral sequence reconstruction and distance averaging. We then give a proof of our main result in the case of homogeneous trees in Section 3. In Section 4, we conclude with a sketch of the proof in the general case.

In the Appendix, we provide a few complimentary results. In Section A, we show that the distance matrix is not in general a sufficient statistic. In Section B, we analyze a standard algorithm, known as WPGMA, in the so-called molecular clock case, that is, when the mutation rate is the same on all branches of the tree. In particular, in the latter case we note that the discretized branch length assumption is not needed.

2 Ancestral Reconstruction and Distance Averaging

Let $\varphi \geq 2$, $0 < \Delta \leq f \leq g < g^* = \ln \sqrt{2}$, and $Q \in \mathbb{Q}_\varphi$ with corresponding stationary distribution $\pi > 0$. In this section we restrict ourselves to the homogeneous case $\mathcal{T} = \mathcal{T}^{(h)} = (V, E, [n], \rho; \tau)$ where we take $h = \log_2 n$ and $f \leq \tau_e \leq g$ and τ_e is an integer multiple of Δ , $\forall e \in E$. (See Examples 1 and 2 and Theorem 2.)⁴

⁴Note that, without loss of generality, we can consider performing ancestral state reconstruction on a homogeneous tree as it is always possible to “complete” a general tree with zero-length edges. We come back to this point in Section 4.

Throughout this section, we use a sequence length $k > \kappa \log(n)$ where κ is a constant to be determined later. We generate k i.i.d. samples $(\xi_V^i)_{i=1}^k$ from the GTR model (\mathcal{T}, Q) with state space Φ .

2.1 Distance Estimator

The standard log-det estimator does not lend itself well to the averaging procedure discussed above. For reconstruction purposes, we instead use an estimator involving the right eigenvector ν corresponding to the second eigenvalue Λ_2 of Q . For $a, b \in [n]$, consider the estimator

$$\hat{\tau}(a, b) = -\ln \left(\nu^\top \hat{F}^{ab} \nu \right), \quad (4)$$

where the correlation matrix \hat{F}^{ab} was introduced in Definition 7. We first give a proof that this is indeed a legitimate distance estimator. For more on connections between eigenvalues of the rate matrix and distance estimation, see e.g. [GL96, GL98, GMY09].

Lemma 1 (Distance Estimator) *Let $\hat{\tau}$ be as above. For all $a, b \in [n]$, we have*

$$\mathbb{E}[e^{-\hat{\tau}(a,b)}] = e^{-\tau(a,b)}.$$

Proof: Note that $\mathbb{E}[\hat{F}_{ij}^{ab}] = \pi_i (e^{-\tau(a,b)Q})_{ij}$. Then

$$\begin{aligned} \mathbb{E} \left[\nu^\top \hat{F}^{ab} \nu \right] &= \sum_{i \in \Phi} \nu_i \sum_{j \in \Phi} \pi_i \left(e^{-\tau(a,b)Q} \right)_{ij} \nu_j \\ &= \sum_{i \in \Phi} \nu_i (\pi_i e^{-\tau(a,b)} \nu_i) \\ &= e^{-\tau(a,b)} \sum_{i \in \Phi} \pi_i \nu_i^2 \\ &= e^{-\tau(a,b)}. \end{aligned}$$

■

For $a \in [n]$ and $i = 1, \dots, k$, let

$$\sigma_a^i = \nu_{\xi_a^i}.$$

Then (4) is equivalent to

$$\hat{\tau}(a, b) = -\ln \left(\frac{1}{k} \sum_{i=1}^k \sigma_a^i \sigma_b^i \right). \quad (5)$$

Note that in the CFN case, we have simply $\nu = (1, -1)^\top$ and hence (5) can be interpreted as a generalization of the CFN metric.

2.2 Ancestral Sequence Reconstruction

Let $e = (x, y) \in E$ and assume that x is closest to ρ (in topological distance). We define $\text{Path}(\rho, e) = \text{Path}(\rho, y)$, $|e|_\rho = |\text{Path}(\rho, e)|$, and

$$R_\rho(e) = (1 - \theta_e^2) \Theta_{\rho, y}^{-2},$$

where $\Theta_{\rho, y} = e^{-\tau(\rho, y)}$ and $\theta_e = e^{-\tau(e)}$.

Proposition 1 below is a variant of Lemma 5.3 in [MP03]. For completeness, we give a proof.

Proposition 1 (Weighted Majority: GTR Version) *Let $\xi_{[n]}$ be a sample from $\mathcal{L}_{\mathcal{T}, Q}$ (see Definition 4) with corresponding $\sigma_{[n]}$. For a unit flow Ψ from ρ to $[n]$, consider the estimator*

$$S = \sum_{x \in [n]} \frac{\Psi(x) \sigma_x}{\Theta_{\rho, x}}.$$

Then, we have

$$\mathbb{E}[S] = 0,$$

$$\mathbb{E}[S \mid \xi_\rho] = \sigma_\rho,$$

and

$$\text{Var}[S] = 1 + K_\Psi,$$

where

$$K_\Psi = \sum_{e \in E} R_\rho(e) \Psi(e)^2.$$

Proof: We follow the proofs of [EKPS00, MP03]. Let \bar{e}_i be the unit vector in direction i . Let $x \in [n]$, then

$$\mathbb{E}[\bar{e}_{\xi_x}^\top \mid \xi_\rho] = \bar{e}_{\xi_\rho}^\top e^{\tau(\rho, x)Q}.$$

Therefore,

$$\mathbb{E}[\sigma_x \mid \xi_\rho] = \bar{e}_{\xi_\rho}^\top e^{\tau(\rho, x)Q} \nu = \sigma_\rho e^{-\tau(\rho, x)},$$

and

$$\mathbb{E}[S \mid \xi_\rho] = \sum_{x \in [n]} \frac{\Psi(x) \sigma_\rho e^{-\tau(\rho, x)}}{\Theta_{\rho, x}} = \sigma_\rho \sum_{x \in [n]} \Psi(x) = \sigma_\rho.$$

In particular,

$$\mathbb{E}[S] = \sum_{i \in \Phi} \pi_i \nu_i = 0.$$

For $x, y \in [n]$, let $x \wedge y$ be the meeting point of the paths between ρ, x, y . We have

$$\begin{aligned}
\mathbb{E}[\sigma_x \sigma_y] &= \sum_{\iota \in \Phi} \mathbb{P}[\xi_{x \wedge y} = \iota] \mathbb{E}[\sigma_x \sigma_y \mid \xi_{x \wedge y} = \iota] \\
&= \sum_{\iota \in \Phi} \pi_\iota \mathbb{E}[\sigma_x \mid \xi_{x \wedge y} = \iota] \mathbb{E}[\sigma_y \mid \xi_{x \wedge y} = \iota] \\
&= \sum_{\iota \in \Phi} \pi_\iota e^{-\tau(x \wedge y, x)} \nu_\iota e^{-\tau(x \wedge y, y)} \nu_\iota \\
&= e^{-\tau(x, y)} \sum_{\iota \in \Phi} \pi_\iota \nu_\iota^2 \\
&= e^{-\tau(x, y)}.
\end{aligned}$$

Then

$$\begin{aligned}
\text{Var}[S] &= \mathbb{E}[S^2] \\
&= \sum_{x, y \in [n]} \frac{\Psi(x) \Psi(y)}{\Theta_{\rho, x} \Theta_{\rho, y}} \mathbb{E}[\sigma_x \sigma_y] \\
&= \sum_{x, y \in [n]} \Psi(x) \Psi(y) e^{2\tau(\rho, x \wedge y)}.
\end{aligned}$$

For $e \in E$, let $e = (e_\uparrow, e_\downarrow)$ where e_\uparrow is the vertex closest to ρ . Then, by a telescoping sum, for $u \in V$

$$\begin{aligned}
\sum_{e \in \text{Path}(\rho, u)} R_\rho(e) &= \sum_{e \in \text{Path}(\rho, u)} e^{2\tau(\rho, e_\downarrow)} - \sum_{e \in \text{Path}(\rho, u)} e^{2\tau(\rho, e_\uparrow)} \\
&= e^{2\tau(\rho, u)} - 1,
\end{aligned}$$

and therefore

$$\begin{aligned}
\mathbb{E}[S^2] &= \sum_{x, y \in [n]} \Psi(x) \Psi(y) e^{2\tau(\rho, x \wedge y)} \\
&= \sum_{x, y \in [n]} \Psi(x) \Psi(y) \left(1 + \sum_{e \in \text{Path}(\rho, x \wedge y)} R_\rho(e) \right) \\
&= 1 + \sum_{e \in E} R_\rho(e) \sum_{x, y \in [n]} \mathbb{1}\{e \in \text{Path}(\rho, x \wedge y)\} \Psi(x) \Psi(y) \\
&= 1 + \sum_{e \in E} R_\rho(e) \Psi(e)^2.
\end{aligned}$$

■

Let Ψ be a unit flow from ρ to $[n]$. We will use the following multiplicative decomposition of Ψ : If $\Psi(x) > 0$, we let

$$\psi(e) = \frac{\Psi(y)}{\Psi(x)},$$

and, if instead $\Psi(x) = 0$, we let $\psi(y) = 0$. Denoting x_\uparrow the immediate ancestor of $x \in V$ and letting $\theta_x = e^{-\tau(x_\uparrow, x)}$, it will be useful to re-write

$$K_\Psi = \sum_{h'=0}^{h-1} \sum_{x \in L_{h'}^{(h)}} (1 - \theta_x^2) \prod_{e \in \text{Path}(\rho, x)} \frac{\psi(e)^2}{\theta_e^2}, \quad (6)$$

and to define the following recursion from the leaves. For $x \in [n]$,

$$K_{x, \Psi} = 0.$$

Then, let $u \in V - [n]$ with children v_1, v_2 with corresponding edges e_1, e_2 and define

$$K_{u, \Psi} = \sum_{\alpha=1,2} ((1 - \theta_{v_\alpha}^2) + K_{v_\alpha, \Psi}) \left(\frac{\psi(e_\alpha)^2}{\theta_{e_\alpha}^2} \right).$$

Note that, from (6), we have $K_{\rho, \Psi} = K_\Psi$.

Because of our use of short sequences, bounds on the variance are not enough for our purposes: We need exponential concentration on our distance estimates. To obtain such concentration, we give bounds on the exponential moment of S . Our proof generalizes a recent argument of Peres and Roch [PR11].

Proposition 2 (Weighted Majority: Exponential Bound) *For $\zeta \in \mathbb{R}$, let*

$$\Gamma^i(\zeta) = \ln \mathbb{E}[\exp(\zeta S) \mid \xi_\rho = \nu_i].$$

Then, there exists $c > 0$ depending only on Q and f such that for all $\zeta \in \mathbb{R}$, we have

$$\Gamma^i(\zeta) \leq \nu_i \zeta + \frac{1}{2} c \zeta^2 K_\Psi.$$

Proof: We prove the claim by induction, moving away from the leaves. We begin with an analytical lemma inspired by the proof of [PR11].

Lemma 2 (Recursion Step) *Let $M = e^{\tau Q}$ with second right eigenvector ν and corresponding eigenvalue $\lambda = e^{-\tau}$ satisfying $\tau \geq f$. Then there is $c > 0$ depending on Q and f such that for all $i \in \Phi$*

$$F(x) \equiv \sum_{j \in \Phi} M_{ij} \exp(\nu_j x) \leq \exp(\lambda \nu_i x + \frac{1}{2} c (1 - \lambda^2) x^2) \equiv G(x), \quad (7)$$

for all $x \in \mathbb{R}$.

Proof: Let $c' = c(1 - \lambda^2)$. Note that

$$F'(x) = \sum_{j \in \Phi} M_{ij} \nu_j \exp(\nu_j x),$$

$$F''(x) = \sum_{j \in \Phi} M_{ij} \nu_j^2 \exp(\nu_j x),$$

$$G'(x) = (\lambda \nu_i + c' x) \exp(\lambda \nu_i x + \frac{1}{2} c' x^2),$$

and

$$G''(x) = ((\lambda \nu_i + c' x)^2 + c') \exp(\lambda \nu_i x + \frac{1}{2} c' x^2).$$

Hence,

$$F(0) = G(0) = 1,$$

$$F'(0) = G'(0) = \lambda \nu_i.$$

Let

$$\bar{\pi} = \min_{\iota} \pi_{\iota},$$

and

$$\bar{\nu} \equiv \max_i |\nu_i| \leq \frac{1}{\sqrt{\bar{\pi}}}.$$

Note that

$$F''(x) \leq \bar{\nu}^2 \exp(\bar{\nu} |x|) \equiv \bar{F}(x),$$

and

$$G''(x) \geq c' \exp(-\bar{\nu} |x| + \frac{1}{2} c' x^2) \equiv \bar{G}(x).$$

Choose $c' = c^* > 0$ such that $\bar{F}(x) < \bar{G}(x)$ for all $x \in \mathbb{R}$. Note in particular that taking

$$c^* > \max \{ 4\bar{\nu}, \bar{\nu}^2 \exp(2\bar{\nu}) \},$$

is enough. Indeed, for $|x| > 1$ we have $c^* > \bar{\nu}^2$ and $\exp(-\bar{\nu}|x| + \frac{1}{2}c^*x^2) > \exp(\bar{\nu}|x|)$ so that $\bar{F}(x) < \bar{G}(x)$. For $|x| \leq 1$, we have

$$\bar{G}(x) > c^* \exp(-\bar{\nu}) > \bar{\nu}^2 \exp(\bar{\nu}) \geq \bar{F}(x).$$

Now choose $c = c^*(1 - e^{-2f})^{-1}$ in (7) (which implies $c' \geq c^*$ by $\tau \geq f$). Then,

$$G''(x) \geq \bar{G}(x) > \bar{F}(x) \geq F''(x),$$

and therefore

$$G(x) \geq F(x),$$

for all $x \in \mathbb{R}$. ■

Going back to the proof of Proposition 2, let $S_x = \sigma_x$ for all $x \in [n]$ and

$$S_u = \sum_{\alpha=1,2} S_{v_\alpha} \frac{\psi(e_\alpha)}{\theta_{e_\alpha}},$$

where $u \in V - [n]$ with children v_1, v_2 with corresponding edges e_1, e_2 . Note that $S_\rho = S$. Let

$$\Gamma_u^i(\zeta) = \ln \mathbb{E}[\exp(\zeta S_u) \mid \xi_u = i].$$

Take $c > 0$ as in Lemma 2. The main claim is clearly true at the leaves, that is, for all $x \in [n]$

$$\begin{aligned} \Gamma_x^i(\zeta) &= \ln \mathbb{E}[\exp(\zeta S_x) \mid \xi_x = i] \\ &= \ln \mathbb{E}[\exp(\zeta \sigma_x) \mid \xi_x = i] \\ &= \nu_i \zeta \\ &\leq \nu_i \zeta + \frac{1}{2} c \zeta^2 K_{x,\Psi}. \end{aligned}$$

For $u \in V - [n]$ as above, we have by the Markov property, induction, and Lemma 2

$$\begin{aligned}
\Gamma_u^i(\zeta) &= \ln \mathbb{E} \left[\exp \left(\zeta \sum_{\alpha=1,2} S_{v_\alpha} \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) \mid \xi_u = i \right] \\
&= \sum_{\alpha=1,2} \ln \mathbb{E} \left[\exp \left(\zeta S_{v_\alpha} \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) \mid \xi_u = i \right] \\
&= \sum_{\alpha=1,2} \ln \left(\sum_{j \in \Phi} M_{ij}^{e_\alpha} \mathbb{E} \left[\exp \left(\zeta S_{v_\alpha} \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) \mid \xi_{v_\alpha} = j \right] \right) \\
&= \sum_{\alpha=1,2} \ln \left(\sum_{j \in \Phi} M_{ij}^{e_\alpha} \exp \left(\Gamma_{v_\alpha}^j \left(\zeta \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) \right) \right) \\
&\leq \sum_{\alpha=1,2} \ln \left(\sum_{j \in \Phi} M_{ij}^{e_\alpha} \exp \left(\nu_j \left(\zeta \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) + \frac{1}{2} c K_{v_\alpha, \Psi} \left(\zeta \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right)^2 \right) \right) \\
&= \frac{1}{2} c \zeta^2 \sum_{\alpha=1,2} K_{v_\alpha, \Psi} \left(\frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right)^2 \\
&\quad + \sum_{\alpha=1,2} \ln \left(\sum_{j \in \Phi} M_{ij}^{e_\alpha} \exp \left(\nu_j \left(\zeta \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) \right) \right) \\
&\leq \frac{1}{2} c \zeta^2 \sum_{\alpha=1,2} K_{v_\alpha, \Psi} \left(\frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right)^2 \\
&\quad + \sum_{\alpha=1,2} \theta_{e_\alpha} \nu_i \left(\zeta \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right) + \frac{1}{2} c (1 - \theta_{v_\alpha}^2) \left(\zeta \frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right)^2 \\
&= \nu_i \zeta + \frac{1}{2} c \zeta^2 \sum_{\alpha=1,2} ((1 - \theta_{v_\alpha}^2) + K_{v_\alpha, \Psi}) \left(\frac{\psi(e_\alpha)}{\theta_{e_\alpha}} \right)^2 \\
&= \nu_i \zeta + \frac{1}{2} c \zeta^2 K_{u, \Psi}.
\end{aligned}$$

■

2.3 Distance Averaging

The input to our tree reconstruction algorithm is the matrix of all estimated distances between pairs of *leaves* $\{\hat{\tau}(a, b)\}_{a, b \in [n]}$. For short sequences, these estimated distances are known to be accurate for leaves that are close enough. We now

show how to compute distances between internal nodes in a way that involves only $\{\hat{\tau}(a, b)\}_{a, b \in [n]}$ (and previously computed internal weights) using Proposition 2.

Let $0 \leq h' < h$. For $v \in L_{h'}^{(h)}$, let $T_v = (V_v, E_v)$ be the subtree of $T = T^{(h)}$ rooted at v with leaf set denoted L_v . Let $a, b \in L_{h'}^{(h)}$. For $x \in \{a, b\}$, denote by X the leaves of $T = T^{(h)}$ below x . Assume that we are given θ_e , for all e below a, b . We estimate $\tau(a, b)$ as follows

$$\bar{\tau}(a, b) \equiv -\ln \left(\frac{1}{|A||B|} \sum_{a' \in A} \sum_{b' \in B} \Theta_{a, a'}^{-1} \Theta_{b, b'}^{-1} e^{-\hat{\tau}(a', b')} \right).$$

This choice of estimator is suggested by the following observation

$$\begin{aligned} e^{-\bar{\tau}(a, b)} &\equiv \sum_{a' \in A} \sum_{b' \in B} 2^{-2h'} \Theta_{a, a'}^{-1} \Theta_{b, b'}^{-1} e^{-\hat{\tau}(a', b')} \\ &= \frac{1}{k} \sum_{i=1}^k \left(\sum_{a' \in A} \frac{2^{-h'} \sigma_{a'}^i}{\Theta_{a, a'}} \right) \left(\sum_{b' \in B} \frac{2^{-h'} \sigma_{b'}^i}{\Theta_{b, b'}} \right). \end{aligned}$$

Note that the first line depends only on estimates $(\hat{\tau}(u, v))_{u, v \in [n]}$ and $\{\Theta_{v, \cdot}\}_{v \in V_a \cup V_b}$. The last line is the empirical distance between the reconstructed states at a and b when the flow is chosen to be homogeneous in Proposition 1.

Lemma 3 (Large Deviations) *Let $0 \leq h' < h$ and let $a, b \in L_{h'}^{(h)}$. For $x = a, b$, let*

$$S_x = \sum_{x' \in X} \frac{2^{-h'} \sigma_{x'}}{\Theta_{x, x'}}.$$

It holds that

$$\mathbb{E}[e^{-\bar{\tau}(a, b)}] = e^{-\tau(a, b)},$$

and there exists $\zeta^ > 0$ small enough such that*

$$\mathbb{E}[\exp(\zeta S_a S_b)] < +\infty,$$

for all $|\zeta| < |\zeta^|$. In particular, for all $\varepsilon > 0$ there exists $0 < \chi < 1$ such that*

$$\mathbb{P} \left[\left| e^{-\bar{\tau}(a, b)} - \mathbb{E}[e^{-\bar{\tau}(a, b)}] \right| > \varepsilon \right] \leq \chi^k.$$

Moreover, χ is a constant independent of h' .

Proof: We first prove the expectation formula. Note that

$$\begin{aligned}
\mathbb{E}[e^{-\bar{\tau}(a,b)}] &= \mathbb{E}\left[\frac{1}{k} \sum_{i=1}^k \left(\sum_{a' \in A} \frac{2^{-h'} \sigma_{a'}^i}{\Theta_{a,a'}}\right) \left(\sum_{b' \in B} \frac{2^{-h'} \sigma_{b'}^i}{\Theta_{b,b'}}\right)\right] \\
&= \mathbb{E}\left[\left(\sum_{a' \in A} \frac{2^{-h'} \sigma_{a'}}{\Theta_{a,a'}}\right) \left(\sum_{b' \in B} \frac{2^{-h'} \sigma_{b'}}{\Theta_{b,b'}}\right)\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\left(\sum_{a' \in A} \frac{2^{-h'} \sigma_{a'}}{\Theta_{a,a'}}\right) \left(\sum_{b' \in B} \frac{2^{-h'} \sigma_{b'}}{\Theta_{b,b'}}\right) \mid \xi_a, \xi_b\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\sum_{a' \in A} \frac{2^{-h'} \sigma_{a'}}{\Theta_{a,a'}} \mid \xi_a\right] \mathbb{E}\left[\sum_{b' \in B} \frac{2^{-h'} \sigma_{b'}}{\Theta_{b,b'}} \mid \xi_b\right]\right] \\
&= \mathbb{E}[\sigma_a \sigma_b] \\
&= e^{-\tau(a,b)},
\end{aligned}$$

where we used that $|A| = |B| = 2^{h'}$.

To prove the large deviation result, it suffices by standard arguments [Dur96] to bound the exponential moment of

$$S_a S_b = \left(\sum_{a' \in A} \frac{2^{-h'} \sigma_{a'}^i}{\Theta_{a,a'}}\right) \left(\sum_{b' \in B} \frac{2^{-h'} \sigma_{b'}^i}{\Theta_{b,b'}}\right).$$

Let N be $\text{Normal}(0, 1)$ and recall that $\mathbb{E}[e^{\zeta N}] = e^{\zeta^2/2}$. By applying Proposition 2 twice and using Fubini's Theorem for positive random variables (see also [PR11]), we get (letting Ψ be the homogeneous flow on T)

$$\begin{aligned}
\mathbb{E}[\exp(\zeta S_a S_b) \mid \xi_a, \xi_b] &\leq \mathbb{E}[\exp(\sigma_a \zeta S_b + \frac{1}{2} c \zeta^2 S_b^2 K_{a,\Psi}) \mid \xi_a, \xi_b] \\
&= \mathbb{E}[\exp(\sigma_a \zeta S_b + \sqrt{c K_{a,\Psi}} \zeta S_b N) \mid \xi_a, \xi_b] \\
&= \mathbb{E}[\exp(S_b(\sigma_a \zeta + \sqrt{c K_{a,\Psi}} \zeta N)) \mid \xi_a, \xi_b] \\
&\leq \mathbb{E}[\exp(\sigma_b(\sigma_a \zeta + \sqrt{c K_{a,\Psi}} \zeta N) \\
&\quad + \frac{1}{2} c(\sigma_a \zeta + \sqrt{c K_{a,\Psi}} \zeta N)^2 K_{b,\Psi}) \mid \xi_a, \xi_b] \\
&< +\infty,
\end{aligned}$$

uniformly in σ_a, σ_b for $|\zeta| > 0$ small enough, where we used $|\sigma_a|, |\sigma_b| \leq \bar{\nu} < +\infty$, Cauchy-Schwarz, and

$$\mathbb{E}[e^{c^2 \zeta^2 K_{a,\Psi} K_{b,\Psi} N^2}] = \left(\frac{1}{1 - 2(c^2 \zeta^2 K_{a,\Psi} K_{b,\Psi})}\right)^{1/2} < +\infty,$$

for small enough ζ . Above we used the moment-generating function of the chi-square distribution.⁵

To prove that the large deviation result is independent of the level h' , we show that $K_{a,\Psi}$ is uniformly bounded in h' . From (6), we have

$$\begin{aligned}
K_{a,\Psi} &\leq \sum_{i=0}^{h'-1} (1 - e^{-2g}) 2^{h'-i} \frac{e^{2(h'-i)g}}{2^{2(h'-i)}} \\
&\leq \sum_{j=1}^{h'} e^{2jg} e^{-(2 \ln \sqrt{2})j} \\
&= \sum_{j=1}^{h'} e^{2j(g-g^*)} \\
&\leq \sum_{j=0}^{+\infty} (e^{-2(g^*-g)})^j \\
&= \frac{1}{1 - e^{-2(g^*-g)}} < +\infty,
\end{aligned} \tag{8}$$

where recall that $g^* = \ln \sqrt{2}$ and $g < g^*$. ■

In the next section, we use the previous lemma in two situations: 1) to estimate the distance between two close vertices on the same level; 2) to detect that two vertices on the same level are “far apart.” These specializations of Lemma 3 are stated below. We only sketch the proofs, which are straightforward.

Proposition 3 (Deep Distance Computation: Small Diameter) *Let $D > 0$, $\gamma > 0$, and $\varepsilon > 0$. Let $a, b \in L_{h'}^{(h)}$ as above. There exist $\kappa > 0$ such that if the following conditions hold:*

- [Small Diameter] $\tau(a, b) < D$,
- [Sequence Length] $k > \kappa \log(n)$,

then

$$|\bar{\tau}(a, b) - \tau(a, b)| < \varepsilon,$$

with probability at least $1 - O(n^{-\gamma})$.

⁵A more careful analysis gives the dependence of χ in Δ as $\chi = 1 - O(\Delta^2)$ [Roc10].

Proof: Let

$$\varepsilon' = \min\{(e^\varepsilon - 1)e^{-D}, (1 - e^{-\varepsilon})e^{-D}\},$$

and observe that

$$\begin{aligned} \bar{\tau}(a, b) - \tau(a, b) &< -\varepsilon \\ \implies e^{-\bar{\tau}(a, b)} &> e^{-\tau(a, b) + \varepsilon} \\ \implies e^{-\bar{\tau}(a, b)} - e^{-\tau(a, b)} &> (e^\varepsilon - 1)e^{-D} \geq \varepsilon'. \end{aligned}$$

A similar implication holds in the other direction. The result now follows from Lemma 3. ■

Proposition 4 (Deep Distance Computation: Diameter Test) *Let $D > 0$, $W > 5$, and $\gamma > 0$. Let $a, b \in L_{h'}^{(h)}$ as above. There exists $\kappa > 0$ such that if the following conditions hold:*

- [Large Diameter] $\tau(a, b) > D + \ln W$,
- [Sequence Length] $k > \kappa \log(n)$,

then

$$\bar{\tau}(a, b) > D + \ln \frac{W}{2},$$

with probability at least $1 - n^{-\gamma}$. On the other hand, if the first condition above is replaced by

- [Small Diameter] $\tau(a, b) < D + \ln \frac{W}{5}$,

then

$$\bar{\tau}(a, b) \leq D + \ln \frac{W}{4},$$

with probability at least $1 - n^{-\gamma}$.

Proof: The proof is similar to the proof of Proposition 3. ■

3 Reconstructing Homogeneous Trees

In this section, we prove our main result in the case of homogeneous trees. More precisely, we prove the following.

Theorem 4 (Main Result: Homogeneous Case) *Let $0 < \Delta \leq f \leq g < +\infty$ and denote by $\mathbb{HY}_{\Delta}^{f,g}$ the set of all homogeneous phylogenies $\mathcal{T} = (V, E, [n], \rho; \tau)$ satisfying $f \leq \tau_e \leq g$ and τ_e is an integer multiple of Δ , $\forall e \in E$. Let $g^* = \ln \sqrt{2}$. Then, for all $\varphi \geq 2$, $0 < \Delta \leq f \leq g < g^*$ and $Q \in \mathbb{Q}_{\varphi}$, there is a distance-based method solving the phylogenetic reconstruction problem on $\mathbb{HY}_{\Delta}^{f,g} \otimes \{Q\}$ with $k = O(\log n)$.*

In the homogeneous case, we can build the tree level by level using simple “four-point” techniques [Bun71]. See e.g. [SS03, Fel04] for background and details. See also Section 3.2 below. The underlying combinatorial algorithm we use here is essentially identical to the one used by Mossel in [Mos04]. From Propositions 3 and 4, we get that the “local metric” on each level is accurate as long as we compute adequate weights. We summarize this fact in the next proposition. For $\Delta > 0$ and $z \in \mathbb{R}_+$, we let $[z]_{\Delta}$ be the closest multiple of Δ to z (breaking ties arbitrarily). For $D > 0$, $W > 5$, we define

$$\overline{\text{SD}}(a, b) = \mathbb{1} \left\{ [\bar{\tau}(a, b)]_{\Delta} \leq D + \ln \frac{W}{3} \right\},$$

and we let

$$\bar{d}(a, b) = \begin{cases} [\bar{\tau}(a, b)]_{\Delta}, & \text{if } \overline{\text{SD}}(a, b) = 1, \\ +\infty, & \text{o.w.} \end{cases}$$

Proposition 5 (Deep Distorted Metric) *Let $D > 0$, $W > 5$, and $\gamma > 0$. Let $\mathcal{T} = (V, E, [n], \rho; \tau) \in \mathbb{HY}_{\Delta}^{f,g}$ with $g < g^*$. Let $a, b \in L_{h'}^{(h)}$ for $0 \leq h' < h$. Assume we are given, for $x = a, b$, θ_e for all $e \in V_x$. There exists $\kappa > 0$, such that if the following condition holds:*

- [Sequence Length] *The sequence length is $k > \kappa \log(n)$,*

then we have, with probability at least $1 - O(n^{-\gamma})$,

$$\bar{d}(a, b) = \tau(a, b)$$

under either of the following two conditions:

1. [Small Diameter] $\tau(a, b) < D$, or
2. [Finite Estimate] $\bar{d}(a, b) < +\infty$.

Proof: We let $\varepsilon < \Delta/2$. The first part of the proposition follows immediately from Proposition 3 and the second part of Proposition 4. For the second part, choose κ so as to satisfy the conditions of Proposition 3 with diameter $D + \ln W$ and apply the first part of Proposition 4. ■

It remains to show how to compute the weights, which is the purpose of the next section.

3.1 Estimating Averaging Weights

Proposition 5 relies on the prior computation of the weights θ_e for all $e \in V_x$, for $x = a, b$. In this section, we show how this estimation is performed.

Let $a, b, c \in L_{h'}^{(h)}$. Denote by z the meeting point of the paths joining a, b, c . We define the “three-point” estimate

$$\hat{\theta}_{z,a} = \mathbb{O}(a; b, c) \equiv \exp \left(-\frac{1}{2} [\bar{d}(a, b) + \bar{d}(a, c) - \bar{d}(b, c)] \right).$$

Note that the expression in parenthesis is an estimate of the distance between a and z .

Proposition 6 (Averaging Weight Estimation) *Let $a, b, c \in L_{h'}^{(h)}$ as above. Assume that the assumptions of Propositions 3, 4, 5 hold. Assume further that the following condition hold:*

- [Small Diameter] $\tau(a, b), \tau(a, c), \tau(b, c) < D + \ln W$,

then

$$\hat{\theta}_{z,a} = \theta_{z,a},$$

with probability at least $1 - O(n^{-\gamma})$ where $\hat{\theta}_{z,a} = \mathbb{O}(a; b, c)$.

Proof: The proof follows immediately from Proposition 5 and the remark above the statement of Proposition 6. ■

3.2 Putting it All Together

Let $0 \leq h' < h$ and $\mathcal{Q} = \{a, b, c, d\} \subseteq L_{h'}^{(h)}$. The topology of $T^{(h)}$ restricted to \mathcal{Q} is completely characterized by a bi-partition or *quartet split* q of the form: $ab|cd$, $ac|bd$ or $ad|bc$. The most basic operation in quartet-based reconstruction algorithms is the inference of such quartet splits. In distance-based methods in particular, this is usually done by performing the so-called *four-point test*: letting

$$\mathcal{F}(ab|cd) = \frac{1}{2} [\tau(a, c) + \tau(b, d) - \tau(a, b) - \tau(c, d)],$$

we have

$$q = \begin{cases} ab|cd & \text{if } \mathcal{F}(a, b|c, d) > 0 \\ ac|bd & \text{if } \mathcal{F}(a, b|c, d) < 0 \\ ad|bc & \text{o.w.} \end{cases}$$

Of course, we cannot compute $\mathcal{F}(a, b|c, d)$ directly unless $h' = 0$. Instead we use Proposition 5.

Deep Four-Point Test. Assume we have previously computed weights θ_e for all $e \in V_x$, for $x = a, b, c, d$. We let

$$\overline{\mathcal{F}}(ab|cd) = \frac{1}{2}[\overline{d}(a, c) + \overline{d}(b, d) - \overline{d}(a, b) - \overline{d}(c, d)], \quad (9)$$

and we define the *deep four-point test*

$$\overline{\mathbb{FP}}(a, b|c, d) = \mathbb{1}\{\overline{\mathcal{F}}(ab|cd) > f/2\},$$

with $\overline{\mathbb{FP}}(a, b|c, d) = 0$ if any of the distances in (9) is infinite. Also, we extend the *diameter test* $\overline{\mathbb{SD}}$ to arbitrary subsets by letting $\overline{\mathbb{SD}}(\mathcal{S}) = 1$ if and only if $\overline{\mathbb{SD}}(x, y) = 1$ for all pairs $x, y \in \mathcal{S}$.

Algorithm. Fix $D > 4g$, $W > 5$, $\gamma > 3$. Choose κ so as to satisfy Propositions 5 and 6. Let \mathcal{Z}_0 be the set of leaves. The algorithm—a standard cherry picking algorithm—is detailed in Figure 1.

Proof of Theorem 4: The proof of Theorem 4 follows from Propositions 5 and 6. Indeed, at each level h' , we are guaranteed by the above to compute a distorted metric with a radius large enough to detect all cherries on the next level using four-point tests. The proof follows by induction. ■

4 Extension to General Trees

It is possible to generalize the previous arguments to general trees, using a combinatorial algorithm of [DMR11], thereby giving a proof of Theorem 3. To apply the algorithm of [DMR11] we need to obtain a generalization of Proposition 5 for disjoint subtrees in “general position.” This is somewhat straightforward and we give a quick sketch in this section.

4.1 Basic Definitions

The algorithm in [DMR11] is called Blindfolded Cherry Picking. We refer the reader to [DMR11] for a full description of the algorithm, which is somewhat involved. It is very similar in spirit to the algorithm introduced in Section 3.2, except for complications due to the non-homogeneity of the tree. The proof in [DMR11] is modular and relies on two main components: a distance-based *combinatorial* argument which remains unchanged in our setting; and a *statistical* argument which we now adapt. The key to the latter is [DMR11, Proposition 4]. Note that [DMR11, Proposition 4] is *not* distance-based as it relies on a complex ancestral reconstruction function—recursive majority. Our main contribution in this section is to show

Algorithm*Input:* Distance estimates $\{\hat{\tau}(a, b)\}_{a, b \in [n]}$;*Output:* Tree;

- For $h' = 1, \dots, h - 1$,

1. **Four-Point Test.** Let

$$\mathcal{R}_{h'} = \{q = ab|cd : \forall a, b, c, d \in \mathcal{Z}_{h'} \text{ distinct such that } \overline{\mathbb{FP}}(q) = 1\}.$$

2. **Cherry Picking.** Identify the cherries in $\mathcal{R}_{h'}$, that is, those pairs of vertices that only appear on the same side of the quartet splits in $\mathcal{R}_{h'}$. Let

$$\mathcal{Z}_{h'+1} = \{a_1^{(h'+1)}, \dots, a_{2^{h-(h'+1)}}^{(h'+1)}\},$$

be the parents of the cherries in $\mathcal{Z}_{h'}$ 3. **Weight Estimation.** For all $z \in \mathcal{Z}_{h'+1}$,

- (a) Let x, y be the children of z . Choose w to be any other vertex in $\mathcal{Z}_{h'}$ with $\overline{\mathbb{SD}}(\{x, y, w\}) = 1$.

- (b) Compute

$$\hat{\theta}_{z,x} = \mathbb{O}(x; y, w).$$

- (c) Repeat the previous step interchanging the role of x and y .

Figure 1: Algorithm.

how this result can be obtained using the techniques of the previous sections—leading to a fully distance-based reconstruction algorithm.

In order to explain the complications due to the non-homogeneity of the tree and state our main result, we first need to borrow a few definitions from [DMR11].

Basic Definitions. Fix $0 < \Delta \leq f \leq g < g^*$ as in Theorem 3. Let $\mathcal{T} = (V, E, [n], \rho; \tau) \in \mathbb{Y}_{\Delta}^{f,g}$ be a phylogeny with underlying tree $T = (V, E)$. In this section, we sometimes refer to the edge set, vertex set and leaf set of a tree T' as $\mathcal{E}(T')$, $\mathcal{V}(T')$, and $\mathcal{L}(T')$ respectively.

Definition 9 (Restricted Subtree) Let $V' \subseteq V$ be a subset of the vertices of T . The subtree of T restricted to V' is the tree T' obtained by 1) keeping only nodes and edges on paths between vertices in V' and 2) by then contracting all paths composed of vertices of degree 2, except the nodes in V' . We sometimes use the notation $T' = T|_{V'}$. See Figure 2 for an example.

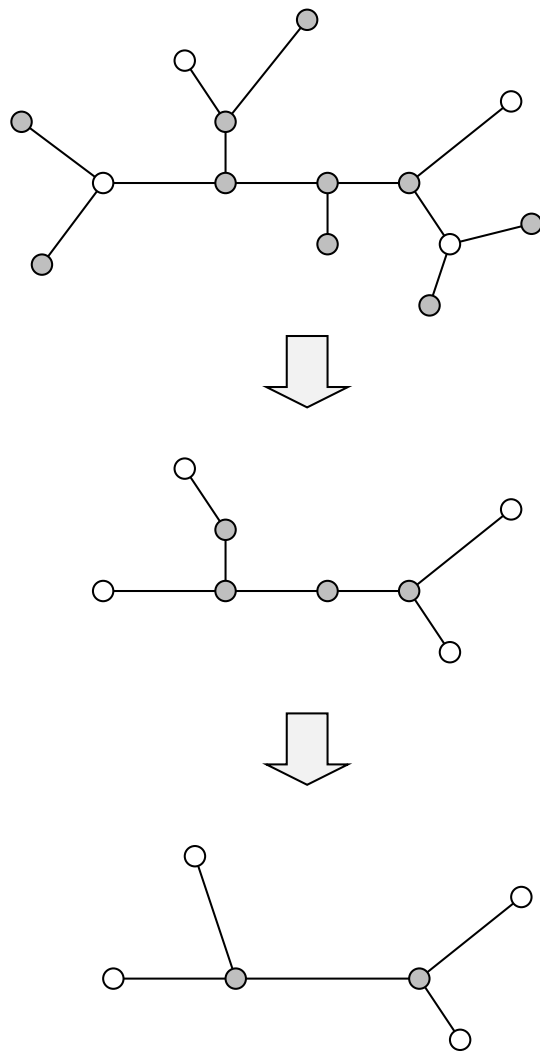


Figure 2: Restricting the top tree to its white nodes.

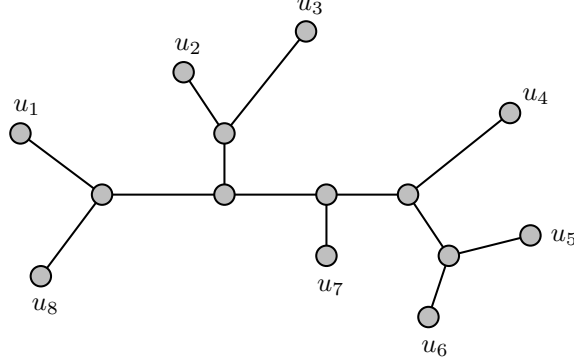


Figure 3: The subtrees $T|_{\{u_1, u_2, u_3, u_8\}}$ and $T|_{\{u_4, u_5, u_6, u_7\}}$ are edge-disjoint. The subtrees $T|_{\{u_1, u_5, u_6, u_8\}}$ and $T|_{\{u_2, u_3, u_4, u_7\}}$ are edge-sharing.

Definition 10 (Edge Disjointness) Denote by $\text{Path}_T(x, y)$ the path (sequence of edges) connecting x to y in T . We say that two restricted subtrees T_1, T_2 of T are edge disjoint if

$$\text{Path}_T(x_1, y_1) \cap \text{Path}_T(x_2, y_2) = \emptyset,$$

for all $x_1, y_1 \in \mathcal{L}(T_1)$ and $x_2, y_2 \in \mathcal{L}(T_2)$. We say that T_1, T_2 are edge sharing if they are not edge disjoint. See Figure 3 for an example.

Definition 11 (Legal Subforest) We say that a tree is a rooted full binary tree if all its internal nodes have degree 3 except the root which has degree 2. A restricted subtree T_1 of T is a legal subtree of T if it is also a rooted full binary tree. We say that a forest

$$\mathcal{F} = \{T_1, T_2, \dots\},$$

is legal subforest of T if the T_i 's are edge-disjoint legal subtrees of T . We denote by $\rho(\mathcal{F})$ the set of roots of \mathcal{F} .

Definition 12 (Dangling Subtrees) We say that two edge-disjoint legal subtrees T_1, T_2 of T are dangling if there is a choice of root for T not in T_1 or T_2 that is consistent with the rooting of both T_1 and T_2 . See Figure 4 below for an example where two legal, edge-disjoint subtrees are not dangling.

Definition 13 (Basic Disjoint Setup (General)) Let $T_1 = T_{x_1}$ and $T_2 = T_{x_2}$ be two restricted subtrees of T rooted at x_1 and x_2 respectively. Assume further that T_1 and T_2 are edge-disjoint, but not necessarily dangling. Denote by y_ℓ, z_ℓ the

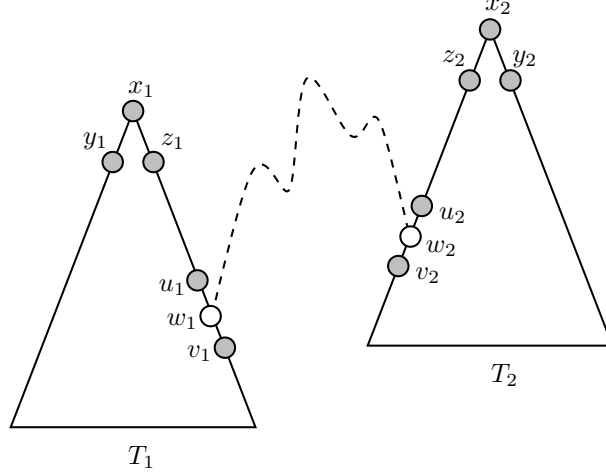


Figure 4: Basic Disjoint Setup (General). The rooted subtrees T_1, T_2 are edge-disjoint but are not assumed to be dangling. The white nodes may not be in the *restricted* subtrees T_1, T_2 . The case $w_1 = x_1$ and/or $w_2 = x_2$ is possible. Note that if we root the tree at any node along the dashed path, the subtrees rooted at y_1 and y_2 are edge-disjoint and dangling (unlike T_1 and T_2).

children of x_ι in T_ι , $\iota = 1, 2$. Let w_ι be the node in T where the path between T_1 and T_2 meets T_ι , $\iota = 1, 2$. Note that w_ι may not be in T_ι since T_ι is restricted, $\iota = 1, 2$. If $w_\iota \neq x_\iota$, assume without loss of generality that w_ι is in the subtree of T rooted at z_ι , $\iota = 1, 2$. We call this configuration the Basic Disjoint Setup (General). See Figure 4. Let $\tau(T_1, T_2)$ be the length of the path between w_1 and w_2 in the metric τ .

4.2 Deep Distorted Metric

Our reconstruction algorithm for homogeneous trees (see Section 3) builds the tree level by level and only encounters situations where one has to compute the distance between two *dangling* subtrees (that is, the path connecting the subtrees “goes above them”). However, when reconstructing general trees by growing a subforest from the leaves, more general situations such as the one depicted in Figure 4 cannot be avoided and have to be dealt with carefully.

Hence, our goal in this subsection is to compute the distance between the internal nodes x_1 and x_2 in the Basic Disjoint Setup (General). We have already shown

how to perform this computation when T_1 and T_2 are *dangling*, as this case is handled easily by Proposition 5 (after a slight modification of the distance estimate; see below). However, in the general case depicted in Figure 4, there is a complication. When T_1 and T_2 are *not* dangling, the reconstructed sequences at x_1 and x_2 are *not* conditionally independent. But it can be shown that for the algorithm Blindfolded Cherry Picking to work properly, we need: 1) to compute the distance between x_1 and x_2 correctly when the two subtrees are close and dangling; 2) detect when the two subtrees are far apart (but an accurate distance estimate is not required in that case). This turns out to be enough because the algorithm Blindfolded Cherry Picking ensures roughly that close reconstructed subtrees are always dangling. We refer the reader to [DMR11] for details.

The key point is the following: if one computes the distance between y_1 and y_2 *rather than* the distance between x_1 and x_2 , then the dangling assumption is satisfied (re-root the tree at any node along the path connecting w_1 and w_2). However, when the algorithm has only reconstructed T_1 and T_2 , we cannot tell which pair in $\{y_1, z_1\} \times \{y_2, z_2\}$ is the right one to use for the distance estimation. Instead, we compute the distance for all pairs in $\{y_1, z_1\} \times \{y_2, z_2\}$ and the following then holds: in the dangling case, all these distances will agree (after subtracting the length of the edges between x_1, x_2 and $\{y_1, z_1, y_2, z_2\}$); in the general case, at least one is correct. This is the basic observation behind the routine DISTORTED-METRIC in Figure 5 and the proof of Proposition 7 below. We slightly modify the definitions of Section 3.

Using the notation of Definition 13, fix $(a, b) \in \{y_1, z_1\} \times \{y_2, z_2\}$. For $x = a, b$, denote by X the leaves of T_x and let $|\ell|_x$ be the graph distance (that is, the number of edges) between x and leaf $\ell \in X$. Assume that we are given θ_e for all $e \in \mathcal{E}(T_a) \cup \mathcal{E}(T_b)$. We estimate $\tau(a, b)$ as follows

$$\bar{\tau}(a, b) \equiv -\ln \left(\sum_{a' \in A} \sum_{b' \in B} 2^{-|a'|_a - |b'|_b} \Theta_{a,a'}^{-1} \Theta_{b,b'}^{-1} e^{-\hat{\tau}(a', b')} \right).$$

Note that, because the tree is binary, it holds that

$$\sum_{a' \in A} \sum_{b' \in B} 2^{-|a'|_a - |b'|_b} = \sum_{a' \in A} 2^{-|a'|_a} \sum_{b' \in B} 2^{-|b'|_b} = 1,$$

and we can think of the weights on A (similarly for B) as resulting from a homogeneous flow Ψ_a from a to A . Then, the bounds on the variance and the exponential moment of

$$S_a \equiv \sum_{a' \in A} 2^{-|a'|_a} \Theta_{a,a'}^{-1} \sigma_{a'},$$

in Propositions 1 and 2 still hold with

$$K_{a, \Psi_a} = \sum_{e \in \mathcal{E}(T_a)} R_a(e) \Psi(e)^2.$$

Moreover K_{a, Ψ_a} is uniformly bounded following an argument identical to (8) in the proof of Lemma 3. In particular, the same large deviations result hold for $\bar{\tau}(a, b)$.

For $D > 0, W > 5$, we define

$$\overline{\text{SD}}(a, b) = \mathbb{1} \left\{ [\bar{\tau}(a, b)]_{\Delta} \leq D + \ln \frac{W}{3} \right\},$$

and we let

$$\bar{d}(a, b) = \begin{cases} [\bar{\tau}(a, b)]_{\Delta}, & \text{if } \overline{\text{SD}}(a, b) = 1, \\ +\infty, & \text{o.w.} \end{cases}$$

Algorithm DISTORTEDMETRIC

Input: Rooted forest $\mathcal{F} = \{T_1, T_2\}$ rooted at vertices x_1, x_2 ; weights τ_e , for all $e \in \mathcal{E}(T_1) \cup \mathcal{E}(T_2)$;

Output: Distance Υ ;

- [Children] Let y_{ι}, z_{ι} be the children of x_{ι} in \mathcal{F} for $\iota = 1, 2$ (if x_{ι} is a leaf, set $z_{\iota} = y_{\iota} = x_{\iota}$);
- [Distance Computations] For all pairs $(a, b) \in \{y_1, z_1\} \times \{y_2, z_2\}$, compute

$$\mathcal{D}(a, b) := \bar{d}(a, b) - \tau(a, x_1) - \tau(b, x_2);$$

- [Multiple Test] If

$$\max \left\{ \left| \mathcal{D}(r_1^{(1)}, r_2^{(1)}) - \mathcal{D}(r_1^{(2)}, r_2^{(2)}) \right| : \right. \\ \left. (r_1^{(\iota)}, r_2^{(\iota)}) \in \{y_1, z_1\} \times \{y_2, z_2\}, \iota = 1, 2 \right\} = 0,$$

return $\Upsilon := \mathcal{D}(z_1, z_2)$, otherwise return $\Upsilon := +\infty$ (return $\Upsilon := +\infty$ if any of the distances above is $+\infty$).

Figure 5: Routine DISTORTEDMETRIC.

Proposition 7 (Accuracy of DISTORTEDMETRIC) *Let $D > 0, W > 5, \gamma > 0$ and $g < g' < g^*$. Consider the Basic Disjoint Setup (General) with $\mathcal{F} = \{T_1, T_2\}$ and $\mathcal{Q} = \{y_1, z_1, y_2, z_2\}$. Assume we are given θ_e for all $e \in \mathcal{E}(T_1) \cup \mathcal{E}(T_2)$. Let Υ denote the output of DISTORTEDMETRIC in Figure 5. There exists $\kappa > 0$, such that if the following condition holds:*

- [Edge Length] *It holds that $\tau(e) \leq g'$, $\forall e \in \mathcal{E}(T_x)$, $x \in \mathcal{Q}^6$;*
- [Sequence Length] *The sequence length is $k > \kappa \log(n)$,*

then we have, with probability at least $1 - O(n^{-\gamma})$,

$$\Upsilon = \tau(x_1, x_2)$$

under either of the following two conditions:

1. [Dangling Case] *T_1 and T_2 are dangling and $\tau(T_1, T_2) < D$, or*
2. [Finite Estimate] $\Upsilon < +\infty$.

Proof: The proof, which is a simple combination of the proof of Proposition 5 and the remarks above the statement of Proposition 7, is left out. ■

Full Algorithm. The rest of the Blindfolded Cherry Picking algorithm is unchanged except for an additional step to compute averaging weights as in the algorithm of Section 3. This concludes our sketch of the proof of Theorem 3.

⁶For technical reasons explained in [DMR11], we allow edges slightly longer than the upper bound g .

Acknowledgments

This work was triggered by a discussion with Elchanan Mossel about lower bounds for distance methods, following a talk of Joseph Felsenstein. Elchanan pointed out that the distance matrix has a potentially useful correlation structure. I am also indebted to Yuval Peres [PR11].

References

- [Att99] K. Atteson. The performance of neighbor-joining methods of phylogenetic reconstruction. *Algorithmica*, 25(2-3):251–278, 1999.
- [BH87] Daniel Barry and J. A. Hartigan. Statistical analysis of hominoid molecular evolution. *Statist. Sci.*, 2(2):191–210, 1987. With comments by Stephen Portnoy and Joseph Felsenstein and a reply by the authors.
- [Bun71] P. Buneman. The recovery of trees from measures of dissimilarity. In *Mathematics in the Archaeological and Historical Sciences*, pages 187–395. Edinburgh University Press, Edinburgh, 1971.
- [CH91] Joseph T. Chang and John A. Hartigan. Reconstruction of evolutionary trees from pairwise distributions on current species, 1991.
- [Cha96] Joseph T. Chang. Full reconstruction of Markov models on evolutionary trees: identifiability and consistency. *Math. Biosci.*, 137(1):51–73, 1996.
- [CK01] Miklós Csurös and Ming-Yang Kao. Provably fast and accurate recovery of evolutionary trees through harmonic greedy triplets. *SIAM Journal on Computing*, 31(1):306–322, 2001.
- [Csu02] M. Csurös. Fast recovery of evolutionary trees with thousands of nodes. *J. Comput. Biol.*, 9(2):277–97, 2002.
- [CT06] Benny Chor and Tamir Tuller. Finding a maximum likelihood tree is hard. *J. ACM*, 53(5):722–744, 2006.
- [Day87] William H. E. Day. Computational complexity of inferring phylogenies from dissimilarity matrices. *Bull. Math. Biol.*, 49(4):461–467, 1987.

- [DHJ⁺06] Constantinos Daskalakis, Cameron Hill, Alexander Jaffe, Radu Mihaescu, Elchanan Mossel, and Satish Rao. Maximal accurate forests from distance matrices. In *RECOMB*, pages 281–295, 2006.
- [DMR06] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Optimal phylogenetic reconstruction. In *STOC’06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 159–168, New York, 2006. ACM.
- [DMR09] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Phylogenies without branch bounds: Contracting the short, pruning the deep. In *RECOMB*, pages 451–465, 2009.
- [DMR11] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Evolutionary trees and the ising model on the bethe lattice: a proof of steel’s conjecture. *Probability Theory and Related Fields*, 149:149–189, 2011. 10.1007/s00440-009-0246-2.
- [DS86] William H. E. Day and David Sankoff. Computational complexity of inferring phylogenies by compatibility. *Syst. Zool.*, 35(2):224–229, 1986.
- [Dur96] Richard Durrett. *Probability: theory and examples*. Duxbury Press, Belmont, CA, second edition, 1996.
- [EKPS00] W. S. Evans, C. Kenyon, Y. Peres, and L. J. Schulman. Broadcasting on trees and the Ising model. *Ann. Appl. Probab.*, 10(2):410–433, 2000.
- [ESSW99a] P. L. Erdős, M. A. Steel, L. A. Székely, and T. A. Warnow. A few logs suffice to build (almost) all trees (part 1). *Random Struct. Algor.*, 14(2):153–184, 1999.
- [ESSW99b] P. L. Erdős, M. A. Steel, L. A. Székely, and T. A. Warnow. A few logs suffice to build (almost) all trees (part 2). *Theor. Comput. Sci.*, 221:77–118, 1999.
- [Fel04] J. Felsenstein. *Inferring Phylogenies*. Sinauer, Sunderland, MA, 2004.
- [Gas97] O. Gascuel. BIO-NJ: An improved version of the NJ algorithm based on a simple model of sequence data. *Mol. Biol. Evol.*, 14(7):685–695, 1997.

- [Geo88] H. O. Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, 1988.
- [GF82] R. L. Graham. and L. R. Foulds. Unlikelihood that minimal phylogenies for a realistic biological study can be constructed in reasonable computational time. *Math. Biosci.*, 60:133–142, 1982.
- [GL96] X Gu and W H Li. A general additive distance with time-reversibility and rate variation among nucleotide sites. *Proceedings of the National Academy of Sciences of the United States of America*, 93(10):4671–4676, 1996.
- [GL98] Xun Gu and Wen-Hsiung Li. Estimation of evolutionary distances under stationary and nonstationary models of nucleotide substitution. *Proceedings of the National Academy of Sciences of the United States of America*, 95(11):5899–5905, 1998.
- [GMS08] Ilan Gronau, Shlomo Moran, and Sagi Snir. Fast and reliable reconstruction of phylogenetic trees with very short edges. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 379–388, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [GMY09] Ilan Gronau, Shlomo Moran, and Irad Yavneh. Towards optimal distance functions for stochastic substitutions models. Preprint, 2009.
- [HNW99] D. H. Huson, S. H. Nettles, and T. J. Warnow. Disk-covering, a fast-converging method for phylogenetic tree reconstruction. *J. Comput. Biol.*, 6(3–4), 1999.
- [KZZ03] Valerie King, Li Zhang, and Yunhong Zhou. On the complexity of distance-based evolutionary tree reconstruction. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 444–453, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.
- [Lak94] JA Lake. Reconstructing Evolutionary Trees from DNA and Protein Sequences: Paralinear Distances. *Proceedings of the National Academy of Sciences*, 91(4):1455–1459, 1994.
- [LC06] Michelle R. Lacey and Joseph T. Chang. A signal-to-noise analysis of phylogeny estimation by neighbor-joining: insufficiency of polynomial length sequences. *Math. Biosci.*, 199(2):188–215, 2006.

- [Lig85] Thomas M. Liggett. *Interacting particle systems*, volume 276 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, New York, 1985.
- [LSHP94] PJ Lockhart, MA Steel, MD Hendy, and D Penny. Recovering Evolutionary Trees under a More Realistic Model of Sequence. *Mol Biol Evol*, 11(4):605–612, 1994.
- [MHR09] R. Mihaescu, C. Hill, and S. Rao. Fast phylogeny reconstruction through learning of ancestral sequences. Preprint, 2009.
- [Mos03] E. Mossel. On the impossibility of reconstructing ancestral data and phylogenies. *J. Comput. Biol.*, 10(5):669–678, 2003.
- [Mos04] E. Mossel. Phase transitions in phylogeny. *Trans. Amer. Math. Soc.*, 356(6):2379–2404, 2004.
- [Mos07] E. Mossel. Distorted metrics on trees and phylogenetic forests. *IEEE/ACM Trans. Comput. Bio. Bioinform.*, 4(1):108–116, 2007.
- [MP03] E. Mossel and Y. Peres. Information flow on trees. *Ann. Appl. Probab.*, 13(3):817–844, 2003.
- [MV05] Elchanan Mossel and Eric Vigoda. Phylogenetic MCMC Algorithms Are Misleading on Mixtures of Trees. *Science*, 309(5744):2207–2209, 2005.
- [PR11] Yuval Peres and Sébastien Roch. Reconstruction on trees: Exponential moment bounds for linear estimators. *Electron. Comm. Probab.*, 16:251–261 (electronic), 2011.
- [Roc06] Sébastien Roch. A short proof that phylogenetic tree reconstruction by maximum likelihood is hard. *IEEE/ACM Trans. Comput. Biology Bioinform.*, 3(1):92–94, 2006.
- [Roc08] Sébastien Roch. Sequence-length requirement for distance-based phylogeny reconstruction: Breaking the polynomial barrier. In *FOCS*, pages 729–738, 2008.
- [Roc10] Sebastien Roch. Toward extracting all phylogenetic information from matrices of evolutionary distances. *Science*, 327(5971):1376–1379, 2010.

- [RS96] A Rzhetsky and T Sitnikova. When is it safe to use an oversimplified substitution model in tree- making? *Mol Biol Evol*, 13(9):1255–1265, 1996.
- [SHP88] M. A. Steel, M. D. Hendy, and D. Penny. Loss of information in genetic distances. *Nature*, 336(6195):118, 1988.
- [SN87] N. Saitou and M. Nei. The neighbor-joining method: A new method for reconstructing phylogenetic trees. *Mol. Biol. Evol.*, 4(4):406–425, 1987.
- [SS63] R. Sokal and P. Sneath. *Principles of Numerical Taxonomy*. W. H. Freeman and Co., San Francisco, Calif., 1963.
- [SS99] Michael A. Steel and László A. Székely. Inverting random functions. *Ann. Comb.*, 3(1):103–113, 1999. Combinatorics and biology (Los Alamos, NM, 1998).
- [SS02] M. A. Steel and L. A. Székely. Inverting random functions. II. Explicit bounds for discrete maximum likelihood estimation, with applications. *SIAM J. Discrete Math.*, 15(4):562–575 (electronic), 2002.
- [SS03] C. Semple and M. Steel. *Phylogenetics*, volume 22 of *Mathematics and its Applications series*. Oxford University Press, 2003.
- [Ste94] M. Steel. Recovering a tree from the leaf colourations it generates under a Markov model. *Appl. Math. Lett.*, 7(2):19–23, 1994.
- [Ste01] M. Steel. My Favourite Conjecture. Preprint, 2001.
- [Was04] Larry Wasserman. *All of statistics*. Springer Texts in Statistics. Springer-Verlag, New York, 2004. A concise course in statistical inference.

A The Distance Matrix is Not Sufficient

A statistic (i.e., a function of the full data) is called *sufficient* if, conditioned on the value of the statistic, the distribution of the full data does not depend on the parameters of the generating model. Roughly speaking, a sufficient statistic encapsulates all the information about the data. See e.g. [Was04]. In this section, we show that the pairwise correlation matrices do not constitute a sufficient statistic for the full Markov model of evolution. Hence, there is in principle more information in the full sequence dataset than there is in the matrix of evolutionary distances.

We give a simple example of non-sufficiency. Consider a four-leaf tree with leaf set $L = \{a, b, c, d\}$ and split $ab|cd$. Assume we use a CFN model with purines denoted “0” and pyrimidines denoted “1” with equal mutation probabilities p . Consider the following correlation matrices

$$\widehat{F}_{v_1 v_2}^{ij} = \frac{1}{4},$$

for all $i \neq j \in L$ and $v_1, v_2 \in \{0, 1\}$. Two different datasets consistent with these correlation matrices are

$$\text{Data}_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix},$$

and

$$\text{Data}_2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \end{bmatrix},$$

where the columns are the sites and the rows are the leaves in the order a, b, c, d .

We compare the probability of observing the two datasets under two different values of p : $p = \varepsilon$ and $p = 1/2 - \varepsilon$ for $\varepsilon > 0$ small. In the first case, in a first approximation it suffices to compute the parsimony scores and we have

$$\mathbb{P}_\varepsilon[\text{Data}_1] = \left(\frac{\varepsilon}{2}\right)^8 + O(\varepsilon^9) = \frac{\varepsilon^8}{256} + O(\varepsilon^9),$$

and

$$\mathbb{P}_\varepsilon[\text{Data}_2] = \left(\frac{1}{2}\right)^2 \left(\frac{\varepsilon}{2}\right)^2 (\varepsilon^2)^4 + O(\varepsilon^{11}) = \frac{\varepsilon^{10}}{16} + O(\varepsilon^{11}).$$

In particular, we get the ratio

$$\frac{\mathbb{P}_\varepsilon[\text{Data}_2 \mid \widehat{F}]}{\mathbb{P}_\varepsilon[\text{Data}_1 \mid \widehat{F}]} = \frac{\mathbb{P}_\varepsilon[\text{Data}_2]}{\mathbb{P}_\varepsilon[\text{Data}_1]} = \varepsilon^2 + O(\varepsilon^3).$$

On the other hand, if $p = 1/2 - \varepsilon$ then the state distribution is almost uniform and we get

$$\frac{\mathbb{P}_{1/2-\varepsilon}[\text{Data}_2 \mid \widehat{F}]}{\mathbb{P}_{1/2-\varepsilon}[\text{Data}_1 \mid \widehat{F}]} = \frac{\mathbb{P}_{1/2-\varepsilon}[\text{Data}_2]}{\mathbb{P}_{1/2-\varepsilon}[\text{Data}_1]} = 1 + O(\varepsilon).$$

Since the ratios are different, we have shown that the distribution of the data conditioned on the correlation matrices depends on the parameters of the model. Therefore, the distance matrix is not a sufficient statistic.

B Probabilistic Analysis of WPGMA

Let $0 < f < g < +\infty$ and denote by $\mathbb{U}\mathbb{Y}^{f,g}$ the set of all phylogenies $\mathcal{T} = (V, E, [n], \rho; \tau) \in \mathbb{Y}^{f,g}$ where we have further that τ is ultrametric, that is, for all $v \in V$ it holds that $\tau(v, x) = \tau(v, y) \equiv \tau(v)$, for all leaves x, y below v . This is known as the molecular clock assumption, that is, the case where the mutation rate is equal on all edges. In that case, there are particularly simple clustering algorithms. We recall the WPGMA algorithm in Figure 6. In the molecular clock case, it is enough to consider “uncorrected” distances [RS96]. Therefore, we run WPGMA with the uncorrected distance estimates

$$\hat{\tau}_u(a, b) = \frac{1 - \hat{\omega}(a, b)}{2},$$

where

$$\hat{\omega}(a, b) = \nu^\top \widehat{F}^{ab} \nu,$$

for $a, b \in [n]$. We call a subset of leaves A a *clade* if it corresponds to all leaf descendants of an internal node a^* called the most recent common ancestor (MRCA). For a clade A with MRCA a^* and a leaf $a \in A$, we let $|a|_A = |a|_{a^*}$ and $\Theta_A = \Theta_{a^*, a}$. For disjoint clades A and B , we let

$$\hat{\tau}_u(A, B) = \sum_{a \in A} \sum_{b \in B} 2^{-|a|_A} 2^{-|b|_B} \hat{\tau}_u(a, b) = \frac{1 - \hat{\omega}(A, B)}{2},$$

where

$$\hat{\omega}(A, B) = \sum_{a \in A} \sum_{b \in B} 2^{-|a|_A} 2^{-|b|_B} \hat{\omega}(a, b).$$

We define

$$\omega(a, b) = e^{-\tau(a, b)},$$

and

$$\tau_u(a, b) = \frac{1 - e^{-\tau(a, b)}}{2}.$$

And similarly for $\omega(A, B)$ and $\tau_u(A, B)$.

Throughout this section, we use a sequence length $k > \kappa \log(n)$ where κ is a constant to be determined later.

Algorithm WPGMA

Input: Distance estimates $\{\hat{\tau}_u(a, b)\}_{a, b \in [n]}$;

Output: Tree;

- **Initialization.** Let \mathcal{Z}_0 be the set of leaves as clusters, that is,

$$\mathcal{Z}_0 = \{\{l\} : l \in [n]\},$$

and for all $a, b \in [n]$ let

$$\hat{\tau}_u(\{a\}, \{b\}) = \hat{\tau}_u(a, b).$$

- **Main Loop.** For $i = 1, \dots, n - 1$,

- **Selection Step.** Let

$$(A^*, B^*) \in \arg \min \{\hat{\tau}_u(A, B) : A, B \in \mathcal{Z}_{i-1} \text{ distinct}\}.$$

Merge clusters A^*, B^* to obtain \mathcal{Z}_i .

- **Reduction Step.** For all $C \in \mathcal{Z}_i - \{A^* \cup B^*\}$, compute

$$\hat{\tau}_u(C, A^* \cup B^*) = \frac{1}{2}[\hat{\tau}_u(C, A^*) + \hat{\tau}_u(C, B^*)]. \quad (10)$$

- **Output.** Output tree implied by the successive clusterings $\mathcal{Z}_0, \dots, \mathcal{Z}_{n-1}$.

Figure 6: Algorithm WPGMA.

Theorem 5 (Analysis of WPGMA) *For all $0 < f < g < g^*$, WPGMA solves the phylogenetic reconstruction problem on $\mathbb{U}\mathbb{Y}^{f, g} \otimes \{Q\}$ with $k = O(\log n)$.*

Proof: Fix $\overline{D} > 3g + 2f$, $2g + 2f < \underline{D} < \overline{D}$, and

$$\varepsilon' < \min \left\{ \frac{e^{2f} - 1}{e^{2f} + 1}, \frac{e^{\underline{D}-2g-2f} - 1}{e^{\underline{D}-2g-2f} + 1} \right\}.$$

This choice ensures that

$$e^{2f} \frac{1 - \varepsilon'}{1 + \varepsilon'} > 1,$$

and

$$e^{\underline{D} - 2g - 2f} \frac{1 - \varepsilon'}{1 + \varepsilon'} > 1,$$

which will be needed later. Let

$$\varepsilon = \min\{\varepsilon' e^{-\overline{D}}, \varepsilon' e^{-\underline{D}}\},$$

and let χ be as in Lemma 3 for this choice of ε . Taking κ large enough, assume the conclusion of Lemma 3 holds for all pairs of clades in the tree, an event we denote by (\star) .

By definition, we have

$$\hat{\tau}_u(A, B) \leq \hat{\tau}_u(A', B') \iff \hat{\omega}(A, B) \geq \hat{\omega}(A', B').$$

For convenience, in the rest of the proof we work with $\hat{\omega}$ rather than $\hat{\tau}_u$. If A, B are disjoint clades with respective MRCA a^* and b^* satisfying $\tau(a^*, b^*) < \overline{D}$, we have

$$\begin{aligned} \hat{\omega}(A, B) &< \omega(A, B) + \Theta_A \Theta_B \varepsilon \\ &\leq \Theta_A \Theta_B (e^{-\tau(a^*, b^*)} + \varepsilon' e^{-\overline{D}}) \\ &< \Theta_A \Theta_B (e^{-\tau(a^*, b^*)} + \varepsilon' e^{-\tau(a^*, b^*)}) \\ &= \omega(A, B)(1 + \varepsilon'), \end{aligned}$$

and similarly

$$\hat{\omega}(A, B) > \omega(A, B)(1 - \varepsilon').$$

On the other hand, if $\tau(a^*, b^*) > \underline{D}$, we have

$$\begin{aligned} \hat{\omega}(A, B) &< \omega(A, B) + \Theta_A \Theta_B \varepsilon \\ &\leq \Theta_A \Theta_B (e^{-\tau(a^*, b^*)} + \varepsilon' e^{-\underline{D}}) \\ &< \Theta_A \Theta_B (e^{-\underline{D}} + \varepsilon' e^{-\underline{D}}) \\ &= \Theta_A \Theta_B e^{-\underline{D}}(1 + \varepsilon'). \end{aligned}$$

By (\star) these inequalities hold for all such pairs of clades.

Two clades A, B are sister clades if their MRCA is their immediate ancestor. We use the following convention. Recall that the leaves are denoted $\{1, \dots, n\}$. We let $\min A$ be the smallest label in A . When denoting a pair of sister clades (A, B) ,

we always assume $\min A < \min B$. There are $n-1$ pairs of sister clades. Order the sister pairs by decreasing value of $\hat{\omega}(A, B)$, breaking ties by lexicographic order over $(\min A, \min B)$:

$$(A_1, B_1), \dots, (A_{n-1}, B_{n-1}).$$

We assume that WPGMA uses the same tie-breaking rule. We let $C_i = A_i \cup B_i$.

We prove the following basic claim. For all $i = 1, \dots, n-1$, at Selection Step i we choose $(A^*, B^*) = (A_i, B_i)$. The result then follows. We work by induction. For $i = 0$, there is nothing to prove. Assume the claim holds up to some $1 \leq i < n-1$. We make a series of observations:

1. All the current clusters in \mathcal{Z}_{i-1} are clades. This follows from the induction hypothesis. By the induction hypothesis, we also get that the values $\hat{\tau}_u(A, B)$ computed at the Reduction Steps indeed correspond to our original definition:

$$\hat{\tau}_u(A, B) = \sum_{a \in A} \sum_{b \in B} 2^{-|a|_A} 2^{-|b|_B} \hat{\tau}_u(a, b) = \frac{1 - \hat{\omega}(A, B)}{2},$$

where

$$\hat{\omega}(A, B) = \sum_{a \in A} \sum_{b \in B} 2^{-|a|_A} 2^{-|b|_B} \hat{\omega}(a, b).$$

2. We show that for all $C \in \mathcal{Z}_{i-1}$, we have

$$\omega(A_i, B_i) e^{-2f} < \Theta_C^2 \leq \omega(A_i, B_i) e^{2g+2f}.$$

Let $C \in \mathcal{Z}_{i-1}$ such that $C = A \cup B$ for sister clades A, B . By (\star) , we have that

$$\begin{aligned} \Theta_C^2 &= \omega(A, B) \\ &> \hat{\omega}(A, B) (1 + \varepsilon')^{-1} \\ &> \hat{\omega}(A_i, B_i) (1 + \varepsilon')^{-1} \\ &> \omega(A_i, B_i) \frac{1 - \varepsilon'}{1 + \varepsilon'} \\ &> \omega(A_i, B_i) e^{-2f}. \end{aligned}$$

Conversely, if a clade $C = A \cup B$ with sister clades A, B satisfies

$$\Theta_C^2 = \omega(A, B) > \omega(A_i, B_i) e^{2f}, \quad (11)$$

then

$$\begin{aligned}
\hat{\omega}(A, B) &> (1 - \varepsilon')\omega(A, B) \\
&> (1 - \varepsilon')\omega(A_i, B_i)e^{2f} \\
&> (1 - \varepsilon')\omega(A_i, B_i)\frac{1 + \varepsilon'}{1 - \varepsilon'} \\
&> (1 + \varepsilon')\omega(A_i, B_i) \\
&> \hat{\omega}(A_i, B_i),
\end{aligned} \tag{12}$$

so that C must be included in a cluster of \mathcal{Z}_i by our induction hypothesis. In particular, if two sister clades A, B are such that $\Theta_A^2, \Theta_B^2 > \omega(A_i, B_i)e^{2g+2f}$ then (11) is satisfied, that is, $\omega(A, B) > \omega(A_i, B_i)e^{2f}$. By (12), (A, B) would have been selected in a previous iteration by induction. That implies, for all $C \in \mathcal{Z}_{i-1}$,

$$\Theta_C^2 \leq \omega(A_i, B_i)e^{2g+2f}.$$

3. We claim that $A_i, B_i \in \mathcal{Z}_{i-1}$. Indeed, by the previous paragraph all clades with Θ^2 -value at least $\omega(A_i, B_i)e^{2f}$ have been constructed in a previous iteration. In particular, the clade A_i has been constructed in a previous step as it satisfies

$$\Theta_{A_i}e^{-f} > \Theta_{C_i} = \sqrt{\omega(A_i, B_i)}.$$

The same holds for B_i . Moreover, A_i and B_i being sister clades of each other (and no other clades), they cannot have been selected inside another pair by our induction hypothesis.

4. By construction, (A_i, B_i) is chosen over all other sister clades present in \mathcal{Z}_{i-1} . So it remains to show that (A_i, B_i) is selected over all other pairs. Pairs of clades that are far enough will not be selected. That is, if A, B with MRCA a^*, b^* is such that

$$\tau(a^*, b^*) \geq \underline{D},$$

then

$$\begin{aligned}
\hat{\omega}(A, B) &< \Theta_A \Theta_B e^{-\underline{D}}(1 + \varepsilon') \\
&< \omega(A_i, B_i)e^{2g+2f}e^{-\underline{D}}(1 + \varepsilon') \\
&< \hat{\omega}(A_i, B_i)(1 - \varepsilon')^{-1}e^{2g+2f}e^{-\underline{D}}(1 + \varepsilon') \\
&< \hat{\omega}(A_i, B_i),
\end{aligned}$$

by assumption on ε' .

5. Finally, non-sister clades that are closer than \underline{D} cannot be selected. Indeed, assume by contradiction that (A^*, B^*) is such a pair. Since (A^*, B^*) are not sister clades, at least one of them, say A^* without loss of generality, has an immediate ancestor u that is strictly lower than the MRCA of A^* and B^* . Take C^* to be any clade in \mathcal{Z}_{i-1} below u that is different than A^* . There must be such a clade because otherwise A^* would have been merged with its sister already. The MRCA of A^* and C^* is u . Moreover, we must have

$$\Theta_{A^*}^2 > \omega(A_i, B_i)e^{-2f},$$

and

$$\Theta_{C^*}^2 \leq \omega(A_i, B_i)e^{2g+2f},$$

so that

$$\tau(a^*, c^*) < 2g + g + 2f < 3g + 2f < \overline{D},$$

where a^* and c^* are the MRCA of A^* and C^* respectively. Finally by (\star)

$$\begin{aligned} \hat{\omega}(A^*, C^*) &> \omega(A^*, C^*)(1 - \varepsilon') \\ &> \omega(A^*, B^*)e^{2f}(1 - \varepsilon') \\ &> \hat{\omega}(A^*, B^*)(1 + \varepsilon')^{-1}e^{2f}(1 - \varepsilon') \\ &> \hat{\omega}(A^*, B^*). \end{aligned}$$

This is a contradiction.

■